

Multimode Control Attacks on Elections*

Piotr Faliszewski
 Dept. of Computer Science
 AGH University of Science
 and Technology, Kraków
 Poland

Edith Hemaspaandra
 Dept. of Computer Science
 Rochester Institute of Technology
 Rochester, NY 14623
 USA

Lane A. Hemaspaandra
 Dept. of Computer Science
 University of Rochester
 Rochester, NY 14627
 USA

July 11, 2010

Abstract

In 1992, Bartholdi, Tovey, and Trick opened the study of control attacks on elections—attempts to improve the election outcome by such actions as adding/deleting candidates or voters. That work has led to many results on how algorithms can be used to find attacks on elections and how complexity-theoretic hardness results can be used as shields against attacks. However, all the work in this line has assumed that the attacker employs just a single type of attack. In this paper, we model and study the case in which the attacker launches a multipronged (i.e., multimode) attack. We do so to more realistically capture the richness of real-life settings. For example, an attacker might simultaneously try to suppress some voters, attract new voters into the election, and introduce a spoiler candidate. Our model provides a unified framework for such varied attacks, and by constructing polynomial-time multiprong attack algorithms we prove that for various election systems even such concerted, flexible attacks can be perfectly planned in deterministic polynomial time.

1 Introduction

Elections are a central model for collective decision-making: Actors' (voters') preferences among alternatives (candidates) are input to the election rule and a winner (or winners in the case of ties) is declared by the rule. Bartholdi, Orlin, Tovey, and Trick initiated a line of research whose goal is to protect elections from various attacking actions intended to skew

*Also appears as URCS-TR-2010-960.

the election’s results. Bartholdi, Orlin, Tovey, and Trick’s strategy for achieving this goal was to show that for various election systems and attacking actions, even seeing whether for a given set of votes such an attack is possible is NP-complete. Their papers [BTT89a, BO91, BTT92] consider actions such as voter manipulation (i.e., situations where a voter misrepresents his or her vote to obtain some goal) and various types of election control (i.e., situations where the attacker is capable of modifying the structure of an election, e.g., by adding or deleting either voters or candidates). Since then, many researchers have extended Bartholdi, Orlin, Tovey, and Trick’s work by providing new models, new results, and new perspectives. But to the best of our knowledge, until now no one has considered the situation in which an attacker combines multiple standard attack types into a single attack—let us call that a *multipronged* (or *multimode*) *attack*.

Studying multipronged control is a step in the direction of more realistically modeling real-life scenarios. Certainly, in real-life settings an attacker would not voluntarily limit himself or herself to a single type of attack but rather would use all available means of reaching his or her goal. For example, an attacker interested in some candidate p winning might, at the same time, intimidate p ’s most dangerous competitors so that they would withdraw from the election, and encourage voters who support p to show up to vote. In this paper we study the complexity of such multipronged control attacks.¹

Given a type of multiprong control, we seek to analyze its complexity. In particular, we try to show either that one can compute in polynomial time an optimal attack of that control type, or that even recognizing the existence of an attack is NP-hard. It is particularly interesting to ask about the complexity of a multipronged attack whose components each have efficient algorithms. We are interested in whether such a combined attack (a) becomes computationally hard, or (b) still has a polynomial-time algorithm. Regarding the (a) case, we give an example of a natural election system that displays this behavior. Our paper’s core work studies the (b) case and shows that even attacks having multiple prongs can in many cases be planned with perfect efficiency. Such results yield as immediate consequences all the individual efficient attack algorithms for each prong, and as such allow a more compact presentation of results and more compact proofs. But they go beyond that: They show that the interactions between the prongs can be managed without such cost as to move beyond polynomial time.

Related work. Since the seminal paper of Bartholdi, Tovey, and Trick [BTT92], much research has been dedicated to studying the complexity of control in elections. Bartholdi, Tovey, and Trick [BTT92] considered constructive control only, i.e., scenarios where the goal of the attacker is to ensure some candidate’s victory. Hemaspaandra, Hemaspaandra, and Rothe [HHR07] extended their work to the destructive case, i.e., scenarios in which the goal is to prevent someone from winning.

A central but elusive goal of control research is finding a natural election system (with a polynomial-time winner algorithm) that is resistant to all the standard types of control, i.e., for which all the types of control are NP-hard. Hemaspaandra, Hemaspaandra, and

¹In fact, our framework of multiprong control includes the unpriced bribery of [FHH09a] and can be extended to include manipulation.

Rothe [HHR09] showed that there exist highly resistant artificial election systems. Faliszewski et al. [FHHR09a] then showed that the natural system known as Copeland voting is not too far from the goal mentioned above. And Erdélyi, Rothe, and Nowak [ENR09] then showed a system with even more resistances than Copeland, but in a slightly non-standard voter model (see [BEH⁺] for discussion and [EPR10b,EPR10a,Men10] for some related follow-up work).

Recently, researchers also started focusing on the parameterized complexity of control in elections. Faliszewski et al. [FHHR09a] provided several fixed-parameter tractability results. Betzler and Uhlmann [BU09] and Liu et al. [LFZL09] showed so-called W[1]- and W[2]-hardness results for control under various voting rules. In response to the conference version of the present paper [FHH09b], Liu and Zhu conducted a parameterized-complexity study of control in maximin elections [LZ10].

Going in a somewhat different direction, Meir et al. [MPRZ08] bridged the notions of constructive and destructive control by considering utility functions, and in this model obtained control results for multiwinner elections. In multiwinner elections the goal is to elect a whole group of people (consider, e.g., parliamentary elections) rather than just a single person. Elkind, Faliszewski, and Slinko [EFS10] and Chevaleyre et al. [MLCM10] considered two types of problems related to control by adding candidates for the case where it is not known how the voters would rank the added candidates.

Faliszewski et al. [FHHR09c] and Brandt et al. [BBHH10] have studied control (and manipulation and bribery) in so-called single-peaked domains, a model of overall electorate behavior from political science.

There is a growing body of work on manipulation that regards frequency of (non)hardness of election problems (see, e.g., [CS06,FKN08,DP08,XC08b,XC08a,Wal09]). This work studies whether a given NP-hard election problem (to date only manipulation/winner problems have been studied, not control problems) can be often solved in practice (assuming some distribution of votes). Such results are of course very relevant when one's goal is to protect elections from manipulative actions. However, in this paper we typically take the role of an attacker and design control algorithms that are fast on *all* instances.

Faliszewski et al. [FHHR09b,FHH] provide an overview of some complexity-of-election issues.

Organization. In Section 2 we present the standard model of elections and describe relevant voting systems. In Section 3 we introduce multiprong control, provide initial results, and show how existing immunity, vulnerability, and resistance results interact with this model. In Section 4 we provide a complexity analysis of voter and candidate control in maximin elections, showing how multiprong control is useful in doing so. We also show that maximin has an interesting relation to Dodgson elections: No candidate whose Dodgson score is more than m^2 times that of the Dodgson winner(s) can be a maximin winner. In Section 5 we consider fixed-parameter complexity of multiprong control, using as our parameter the number of candidates. Section 6 provides conclusions and open problems.

2 Preliminaries

Elections. An election is a pair (C, V) , where $C = \{c_1, \dots, c_m\}$ is the set of candidates and $V = (v_1, \dots, v_n)$ is a collection of voters. Each voter v_i is represented by his or her preference list.² For example, if we have three candidates, c_1 , c_2 , and c_3 , a voter who likes c_1 most, then c_2 , and then c_3 would have preference list $c_1 > c_2 > c_3$.³ Given an election $E = (C, V)$, by $N_E(c_i, c_j)$, where $c_i, c_j \in C$ and $i \neq j$, we denote the number of voters in V who prefer c_i to c_j . We adopt the following convention for specifying preference lists.

Convention A. *Listing some set D of candidates as an item in a preference list means listing all the members of this set in some fixed, arbitrary order, and listing \overleftarrow{D} means listing all the members of D , but in the reverse order.*

An election system is a mapping that given an election (C, V) outputs a set W , satisfying $W \subseteq C$, called the winners of the election.

We focus on the following five voting systems: plurality, Copeland, maximin, approval, and Condorcet. (However, in Sections 4.4 and 5 we take a detour through some other systems.) Each of plurality, Copeland, maximin, and approval assigns points to candidates and elects those that receive the most points. Let $E = (C, V)$ be an election, where $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$. In plurality, each candidate receives a single point for each voter who ranks him or her first. In maximin, the score of a candidate c_i in E is defined as $\min_{c_j \in C - \{c_i\}} N_E(c_i, c_j)$. For each rational α , $0 \leq \alpha \leq 1$, in Copeland $^\alpha$ candidate c_i receives 1 point for each candidate c_j , $j \neq i$, such that $N_E(c_i, c_j) > N_E(c_j, c_i)$ and α points for each candidate c_j , $j \neq i$, such that $N_E(c_i, c_j) = N_E(c_j, c_i)$. That is, the parameter α describes the value of ties in head-to-head majority contests. In approval, instead of preference lists each voter's ballot is a 0-1 vector, where each entry denotes whether the voter approves of the corresponding candidate (gives the corresponding candidate a point). For example, vector $(1, 0, 0, 1)$ means that the voter approves of the first and fourth candidates, but not the second and third. We use $\text{score}_E(c_i)$ to denote the score of candidate c_i in election E (the particular election system used will always be clear from context).

A candidate c is a Condorcet winner of an election $E = (C, V)$ if for each other candidate $c' \in C$ it holds that $N_E(c, c') > N_E(c', c)$. Clearly, each election has at most one Condorcet winner. (Not every election has a Condorcet winner. However, as our notion of an election allows outcomes in which no one wins, electing the Condorcet winner when there is one and otherwise having no winner is a quite legal election system.)

Computational complexity. We use standard notions of complexity theory, as presented, e.g., in the textbook of Papadimitriou [Pap94]. We assume that the reader is familiar with the complexity classes P and NP, polynomial-time many-one reductions, and the notions of NP-hardness and NP-completeness. \mathbb{N} will denote $\{0, 1, 2, \dots\}$.

²We also assume that each voter has a unique name. However, all the election systems we consider here—except for the election system of Theorem 3.8—are anonymous and thus disregard voter names and the order of the votes.

³Preference lists are also called preference orders, and in this paper we will use these two terms interchangeably.

Most of the NP-hardness proofs in this paper follow by a reduction from the well-known NP-complete problem *exact cover by 3-sets*, known for short as X3C (see, e.g., [GJ79]). In X3C we are given a pair (B, \mathcal{S}) , where $B = \{b_1, \dots, b_{3k}\}$ is a set of $3k$ elements and $\mathcal{S} = \{S_1, \dots, S_n\}$ is a set of 3-subsets of B , and we ask whether there is a subset S' of exactly k elements of \mathcal{S} such that their union is exactly B . We call such a set S' an exact cover of B .

In Section 5, we consider the fixed-parameter complexity of multiprong control. The idea of fixed-parameter complexity is to measure the complexity of a given decision problem with respect to both the instance size (as in the standard complexity theory) and some parameter of the input (in our case, the number of candidates involved). For a problem to be said to be fixed-parameter tractable, i.e., to belong to the complexity class FPT, we as is standard require that the problem can be solved by an algorithm running in time $f(j)n^{O(1)}$, where n is the size of the encoding of the given instance, j is the value of the parameter for this instance, and f is some function. Note that f does not have to be polynomially bounded or even computable. However, in all FPT claims in this paper, f is a computable function. That is, our algorithms actually achieve so-called strongly uniform fixed-parameter tractability. We point readers interested in parameterized complexity to, for example, Niedermeier’s book [Nie06].

3 Control and Multiprong Control

In this section we introduce multiprong control, that is, control types that combine several standard types of control. We first provide the definition, then proceed to analyzing general properties of multiprong control, then consider multiprong control for election systems for which the complexity of single-prong control has already been established, and finally give an example of an election system for which multiprong control becomes harder than any of its constituent prongs (assuming $P \neq NP$).

3.1 The Definition

We consider combinations of control by adding/deleting candidates/voters⁴ and by bribing voters. Traditionally, bribery has not been considered a type of control but it fits the model very naturally and strengthens our results.

In discussing control problems, we must be very clear about whether the goal of the attacker is to make his or her preferred candidate *the only winner*, or is to make his or her preferred candidate *a winner*. To be clear on this, we as is standard will use the term “unique-winner model” for the model in which the goal is to make one’s preferred candidate the one and only winner, and we will use the term “nonunique-winner model” for the approach in which the goal is to make one’s preferred candidate be a winner. (Note that if exactly one person wins, he or she most certainly is considered to have satisfied the

⁴Other control types, defined by Bartholdi, Tovey, and Trick [BT92] and refined by Hemaspaandra, Hemaspaandra, and Rothe [HHR07], regard various types of partitioning candidates and voters.

control action in the nonunique-winner model. The “nonunique” in the model name merely means we are not *requiring* that winners be unique.)

The destructive cases of each of these are, in the nonunique-winner model, blocking one’s despised candidate from being a unique winner,⁵ and in the unique-winner model, blocking one’s despised candidate from being a winner. We take the unique-winner model as the default in this paper, as is the most common model in studies of control.

Definition 3.1. Let \mathcal{E} be an election system. In the unique-winner,⁶ constructive \mathcal{E} -AC+DC+AV+DV+BV control problem we are given:

- (a) two disjoint sets of candidates, C and A ,
- (b) two disjoint collections of voters, V and W , containing voters with preference lists over $C \cup A$,
- (c) a preferred candidate $p \in C$, and
- (d) five nonnegative integers, k_{AC} , k_{DC} , k_{AV} , k_{DV} , and k_{BV} .

We ask whether it is possible to find two sets, $A' \subseteq A$ and $C' \subset C$, and two subcollections of voters, $V' \subseteq V$ and $W' \subseteq W$, such that:

- (e) it is possible to ensure that p is a unique winner of \mathcal{E} election $((C - C') \cup A', (V - V') \cup W')$ via changing preference orders of (i.e., bribing) at most k_{BV} voters in $(V - V') \cup W'$,
- (f) $p \notin C'$, and
- (g) $\|A'\| \leq k_{AC}$, $\|C'\| \leq k_{DC}$, $\|W'\| \leq k_{AV}$, and $\|V'\| \leq k_{DV}$.

In the unique-winner, destructive variant of the problem, we replace item (e) above with: “it is possible to ensure that p is not a unique winner of \mathcal{E} election $((C - C') \cup A', (V - V') \cup W')$ via changing preference orders of at most k_{BV} voters in $(V - V') \cup W'$. (In addition, in the destructive variant we refer to p as “the despised candidate” rather than as “the preferred candidate,” and we often denote him or her by d .)

The phrase AC+DC+AV+DV+BV in the problem name corresponds to four of the standard types of control: adding candidates (AC), deleting candidates (DC), adding voters (AV), deleting voters (DV), and to (unpriced) bribery (BV); we will refer to these five types of control as the *basic* types of control. We again remind the reader that traditionally bribery is not a type of control but we will call it a basic type of control for the sake of uniformity and throughout the rest of the paper we will consider it as such.

⁵We will often use the phrase “a unique winner,” as we just did. The reason we write “a unique winner” rather than “the unique winner” is to avoid the impression that the election necessarily has some (unique) winner.

⁶One can easily adapt the definition to the nonunique-winner model.

Instead of considering all of AC, DC, AV, DV, and BV, we often are interested in some subset of them and so we consider special cases of the AC+DC+AV+DV+BV problem. For example, we write DC+AV to refer to a variant of the AC+DC+AV+DV+BV problem where only deleting candidates and adding voters is allowed. As part of our model we assume that in each such variant, only the parameters relevant to the prongs are part of the input. So, for example, DC+AV would have k_{DC} , k_{AV} , C , V , W , and p as the (only) parts of its input. And the “missing” parts (e.g., for DC+AV, the missing parts are A , k_{AC} , k_{DV} , and k_{BV}) are treated in the obvious way in evaluating the formulas in Definition 3.1, namely, missing sets are treated as \emptyset and missing constants are treated as 0. If we name only a single type of control, we in effect degenerate to one of the standard control problems. We for historical reasons consider also a special case of the AC control type, denoted AC_u (and called control by adding an unlimited number of candidates), where there is no limit on the number of candidates to add, i.e., $k_{AC} = \|A\|$.

There is at least one more way in which we could define multiprong control. The model in the above definition can be called the *separate-resource model*, as the extent to which we can use each basic type of control is bounded separately. In the *shared-resource model* one pool of action allowances must be allocated among the allowed control types (so in the definition above we would replace k_{AC}, k_{DC}, k_{AV} , and k_{DV} with a single value, k , and require that $\|C'\| + \|D'\| + \|V'\| + \|W'\| + \text{the-number-of-bribed-voters} \leq k$). Although one could make various arguments about which model is more appropriate, their computational complexity is related.

Theorem 3.2. *If there is a polynomial-time algorithm for a given variant of multiprong control in the separate-resource model then there is one for the shared-resource model as well.*

Proof. Let \mathcal{E} be an election system. We will describe the idea of our proof on the example of the constructive \mathcal{E} -AC+AV problem. The idea easily generalizes to any other set of allowed control actions (complexity-theory savvy readers will quickly see that we, in essence, give a disjunctive truth-table reduction).

We are given an instance I of the constructive \mathcal{E} -AC+AV problem in the shared-resource model, where k is the limit on the sum of the number of candidates and voters that we may add. Given a polynomial-time algorithm for the separate-resource variant of the problem, we solve I using the following method. (If $k > \|A\| + \|W\|$ then set $k = \|A\| + \|W\|$.) We form a sequence I_0, \dots, I_k of instances of the separate-resource variant of the problem, where each I_ℓ , $0 \leq \ell \leq k$, is identical to I , except that we are allowed to add at most ℓ candidates and at most $k - \ell$ voters. We accept if at least one of I_ℓ is a “yes”-instance of the separate-resource, constructive \mathcal{E} -AC+AV problem. Clearly, this algorithm is correct and runs in polynomial time. \square

It would be interesting to consider a variant of the shared-resource model where various actions come at different costs (e.g., adding some candidate c' might be much more expensive—or difficult—than adding some other candidate c''). This approach would be

close in spirit to priced bribery of [FHH09a]. Analysis of such priced control is beyond the scope of the current paper.

3.2 Susceptibility, Immunity, Vulnerability, and Resistance

As is standard in the election-control (and election-bribery) literature, we consider vulnerability, immunity, susceptibility, and resistance to control. Let \mathcal{E} be an election system and let \mathcal{C} be a type of control. We say that \mathcal{E} is susceptible to constructive (destructive) \mathcal{C} control if there is a scenario in which effectuating \mathcal{C} makes someone become (stop being) a unique winner of some \mathcal{E} election E . \mathcal{E} is immune to constructive (destructive) \mathcal{C} control if \mathcal{E} is not susceptible to constructive (destructive) \mathcal{C} control. We say that \mathcal{E} is vulnerable to constructive (destructive) \mathcal{C} control if \mathcal{E} is susceptible to constructive (destructive) \mathcal{C} and there is a polynomial-time algorithm that decides the constructive (destructive) \mathcal{E} - \mathcal{C} problem. Actually, this paper’s vulnerability algorithms/proofs will each go further and will in polynomial time produce, or will make implicitly clear how to produce, the successful control action. So we in each case are even achieving the so-called certifiable vulnerability of Hemaspaandra, Hemaspaandra, and Rothe [HHR07]. \mathcal{E} is resistant to constructive (destructive) \mathcal{C} control if \mathcal{E} is susceptible to (destructive) \mathcal{C} control and the constructive (destructive) \mathcal{E} - \mathcal{C} problem is NP-hard.

The next three theorems describe how multiprong control problems can inherit susceptibility, immunity, vulnerability, and resistance from the basic control types that they are built from.

Theorem 3.3. *Let \mathcal{E} be an election system and let $C_1 + \dots + C_k$ be a variant of multiprong control (so $1 \leq k \leq 5$ and each C_i is a basic type of control). \mathcal{E} is susceptible to constructive (destructive) $C_1 + \dots + C_k$ control if and only if \mathcal{E} is susceptible to at least one of constructive (destructive) C_1, \dots, C_k control.*

Proof. The “if” direction is trivial: The attacker can always choose to use only the type of control to which \mathcal{E} is susceptible. As to the “only if” direction, it is not hard to see that if there is some input election for which by a $C_1 + \dots + C_k$ action we can achieve our desired change (of creating or removing unique-winnerhood for p , depending on the case), then there is *some* election (not necessarily our input election) for which one of those actions alone achieves our desired change. In essence, we can view a control action A of type $C_1 + \dots + C_k$ as a sequence of operations, each one of one of the C_1, \dots, C_k types, that—when executed in order—transform our input election into an election where our goal is satisfied. Thus there is a single operation within A —and this operation is of one of the types C_1, \dots, C_k —that transforms some election E' where our goal is not satisfied to some election E'' where the goal is satisfied. \square

In the next theorem we show that if a given election system is vulnerable to some basic type of control and is immune to another basic type of control, then it is vulnerable to these two types of control combined. The proof of this theorem is easy, but we need to be particularly careful as vulnerabilities and immunities can behave quite unexpectedly. For

example, it seems that we can assume that if an election system is vulnerable to AV and DV then it should also be vulnerable to BV, because bribing a particular voter can be viewed as first deleting this voter and then adding—in his or her place—a voter with the preference order as required by the briber. (This assumes we have such a voter among the voters we can add, but when arguing susceptibility/immunity we can make this assumption.) However, there is an easy election system that is vulnerable to both AV and DV control, but that is immune to BV control. This system simply says that in an election $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$, the winner is the candidate c_i such that $n \equiv i - 1 \pmod{m}$.⁷

Theorem 3.4. *Let \mathcal{E} be an election system and let $C_1 + \dots + C_k + D_1 + \dots + D_\ell$ be a variant of multiprong control (so $1 \leq k, \ell \leq 5$ and each C_i and each D_i is a basic control type) such that \mathcal{E} is vulnerable to constructive (destructive) $C_1 + \dots + C_k$ control but is immune to constructive (destructive) $D_1 + \dots + D_\ell$ control. \mathcal{E} is vulnerable to $C_1 + \dots + C_k + D_1 + \dots + D_\ell$ control.*

Proof. We will give a proof for the constructive case only. The proof for the destructive case is analogous. Let \mathcal{E} be an election system as in the statement of the theorem and let I be an instance of constructive \mathcal{E} - $C_1 + \dots + C_k + D_1 + \dots + D_\ell$ control, which contains election $E = (C, V)$, information about the specifics of control actions we can implement, and where the goal is to ensure that candidate p is a unique winner. Let us first consider the case where BV is not among $C_1, \dots, C_k, D_1, \dots, D_\ell$.

Let us assume that there is a sequence A of control actions of types $C_1, \dots, C_k, D_1, \dots, D_\ell$, such that (a) applying A to E is a legal control action within I , and (b) applying A to E results in an election E_{C+D} where p is the unique winner. (We take A to be an empty sequence if p is a unique winner of E .) We split the sequence A into a subsequence A_C that contains exactly the actions of types C_1, \dots, C_k , and a subsequence A_D that contains exactly the actions of types D_1, \dots, D_ℓ . Since BV is not among our control actions, it is easy to see that it is possible to apply actions A_C to election E to obtain some election E_C . (To see why it is important that we do not consider BV, assume that BV is among control types C_1, \dots, C_k and AV is among control types D_1, \dots, D_ℓ . In this case, A_C might include an action that bribes a voter that is added by an action from A_D .)

We claim that p is a unique winner of E_C . For the sake of contradiction, let us assume that this is not the case (note that this implies that p is not a unique winner of E). If we apply control actions A_D to E_C , we reach exactly election E_{C+D} , where p is the unique winner. Yet, this is a contradiction, because we assumed that \mathcal{E} is immune to $D_1 + \dots + D_\ell$, i.e., that there is no scenario where control action of type $D_1 + \dots + D_\ell$ makes some candidate a unique winner if he or she was not a unique winner before.

Thus, it is possible to ensure that p is a unique winner by actions of type $C_1 + \dots + C_k$ alone. We chose I arbitrarily, and thus any instance of \mathcal{E} - $C_1 + \dots + C_k + D_1 + \dots + D_\ell$ control can be solved by an algorithm that considers control actions of type $C_1 + \dots + C_k$

⁷Of course, this election system is not neutral; permuting the names of the candidates can change the outcome of an election.

only. This proves that \mathcal{E} is vulnerable to $C_1 + \dots + C_k + D_1 + \dots + D_\ell$ control because, as we have assumed, it is vulnerable to $C_1 + \dots + C_k$ control.

It remains to prove the theorem for the case where BV is among our control actions. In the case where BV is among the control actions but AV is not, or if AV and BV are in the same group of actions (i.e., either both are among the C_i 's or both are among the D_i 's), it is easy to see that the above proof still works. Similarly, if BV is among the D_i 's and AV is among the C_i 's, the above proof works as well. The only remaining case is if our allowed control types include both BV and AV , where BV is among the C_i 's and AV is among the D_i 's.

In this last case, the proof also follows the general structure of the previous construction, except that we have to take care of one issue: It is possible that sequence A_C includes bribery of voters that are to be added by actions from A_D . (We use the same notation as in the main construction.) Let V_{BV} be the collection of voters that A_C requires to bribe, but that are added in A_D . We form a sequence A'_C that is identical to A_C , except that it starts by adding the voters from V_{BV} , and we let A'_D be identical to A_D , except that it no longer includes adding the voters from V_{BV} . Using sequences A'_C and A'_D instead of A_C and A_D , it is easy to show the following: If it is possible to ensure that p is a unique winner in instance I by a legal action of type $C_1 + \dots + C_k + D_1 + \dots + D_\ell$, then it is also possible to do so by a legal action of type $C_1 + \dots + C_k + AV$, where each added voter is also bribed. Thus, given an instance I of $\mathcal{E}\text{-}C_1 + \dots + C_k + D_1 + \dots + D_\ell$ we can solve it using the following algorithm. Let W be the collection of voters that can be added within I and let k_{AV} be the limit on the number of voters that we can add.

1. Let t be $\max(k_{AV}, \|W\|)$.
2. For each i in $\{0, 1, \dots, t\}$ execute the next two substeps.
 - (a) Form instance I' that is identical to I , except i (arbitrarily chosen) voters from W are added to the election.
 - (b) Run the $\mathcal{E}\text{-}C_1 + \dots + C_k$ algorithm on instance I' and accept if it does.
3. If the algorithm has not accepted yet, reject.

It is easy to see that this algorithm is correct and, since \mathcal{E} is vulnerable to $C_1 + \dots + C_k$, works in polynomial time. This completes the proof of the theorem. \square

Theorem 3.5. *Let \mathcal{E} be an election system and let $C_1 + \dots + C_k$, $1 \leq k \leq 5$, be a variant of multiprong control. If for some i , $1 \leq i \leq k$, \mathcal{E} is resistant to constructive (destructive) C_i control, then \mathcal{E} is resistant to constructive (destructive) $C_1 + \dots + C_k$ control.*

Proof. Let C_i be the control type to which \mathcal{E} is resistant. Since \mathcal{E} is susceptible to constructive (destructive) C_i control, it follows by Theorem 3.3 that \mathcal{E} is susceptible to constructive (destructive) $C_1 + \dots + C_k$ control. And since the $\mathcal{E}\text{-}C_i$ constructive (destructive) control problem is essentially (give or take syntax) an embedded subproblem of the $\mathcal{E}\text{-}C_1 + \dots + C_k$ control problem, clearly \mathcal{E} is resistant to $C_1 + \dots + C_k$ control. \square

By combining the above three theorems, we obtain a simple tool that allows us to classify a large number of multiprong control problems based on the properties of their prongs.

Corollary 3.6. *Let \mathcal{E} be an election system and let $C_1 + \dots + C_k$, $1 \leq k \leq 5$, be a variant of multiprong control, such that for each C_i , $1 \leq i \leq k$, \mathcal{E} is resistant, vulnerable, or immune to constructive (destructive) C_i control. If there is an i , $1 \leq i \leq k$, such that \mathcal{E} is resistant to constructive (destructive) C_i control then \mathcal{E} is resistant to constructive (destructive) $C_1 + \dots + C_k$ control. Otherwise, if there is an i , $1 \leq i \leq k$, such that \mathcal{E} is vulnerable to constructive (destructive) C_i control then \mathcal{E} is vulnerable to constructive (destructive) $C_1 + \dots + C_k$ control. Otherwise, \mathcal{E} is immune to constructive (destructive) $C_1 + \dots + C_k$ control.*

Theorem 3.5 immediately yields many “free” resistance results based on the previous work on control. However, we will focus on the more interesting issue of proving that even multiprong control is easy for some election systems whose control has already been studied (in Section 3.3) and for candidate control in maximin (Section 4).

In general, we do not consider partition cases of control in this paper. However, we make an exception for the next example, which shows how even types of control to which a given election system is immune may prove useful in multiprong control. In constructive control by partition of candidates (reminder: this is not a basic control type) in the ties-eliminate model (PC-TE control type), we are given an election $E = (C, V)$ and a preferred candidate $p \in C$, and we ask whether it is possible to find a partition (C_1, C_2) of C (i.e., $C_1 \cup C_2 = C$ and $C_1 \cap C_2 = \emptyset$) such that p is a winner of the following two-round election: We first find the winner sets, W_1 and W_2 , of elections (C_1, V) and (C_2, V) . If W_1 (W_2) contains more than one candidate, we set $W_1 = \emptyset$ ($W_2 = \emptyset$), since we are in the “ties eliminate” model. The candidates who win election $(W_1 \cup W_2, V)$ are the winners of the overall two-stage election.

Now, let us look at constructive approval-AC+PC-TE control, where (by definition, let us say) we first add new candidates and then perform the partition action. We consider an approval election with two candidates, p and c , where p has 50 approvals and c has 100. We are also allowed to add candidate c' , who has 100 approvals. Clearly, it is impossible to make p a unique winner by adding c' . Exercising the partition action alone does not ensure p 's victory either. However, combining both AC and PC-TE does the job. If we first add c' to the election and then partition candidates into $\{p\}$ and $\{c, c'\}$ then, due to the ties-eliminate rule, p becomes the unique winner. It is rather interesting that even though approval is immune to constructive AC control, there are cases where one has to apply AC control to open the possibility of effectively using other types of control.

The above example is perhaps surprising in light of Theorem 3.4. In essence, in the proof of that theorem we argue that if an election system is vulnerable to some basic control type C but is immune to some other basic control type D , then it is also vulnerable to control type $C + D$. We proved the theorem by showing that we can safely disregard the actions of type D (assuming C does not include BV control type). The above example shows that this proof approach would not work if we considered PC-TE in addition to the basic control types.

3.3 Combining Vulnerabilities

In the previous section we considered the case where separate prongs of a multiprong control problem have different computational properties, e.g., some are resistant, some are vulnerable, and some are immune. In this section we consider the case where an election system is vulnerable to each prong separately, and we show how such vulnerabilities combine within election systems for which control results were obtained in previous papers (see Table 2). In particular, in the next theorem we show that for all the election systems considered in [BTT92], [HHR07], and [FHHR09a], all constructive vulnerabilities to AC, DC, AV, DV, and BV combine to vulnerabilities, and all destructive vulnerabilities to AC, DC, AV, DV, BV combine to vulnerabilities.⁸ That is, for each election system studied in these three papers, if it is separately vulnerable to some basic control types C_1, \dots, C_k , where each $C_i \in \{\text{AC}, \text{DC}, \text{AV}, \text{DV}, \text{BV}\}$, it is also vulnerable to $C_1 + \dots + C_k$.

Theorem 3.7. (a) *Plurality is vulnerable to both constructive AV+DV+BV control and destructive AV+DV+BV control.* (b) *Both Condorcet and approval are vulnerable to AC+AV+DV+BV destructive control.* (c) *For each rational α , $0 \leq \alpha \leq 1$, Copeland $^\alpha$ is vulnerable to destructive AC+DC control.*

Proof. (a) Let us consider an instance I of constructive plurality-AV+DV+BV control where we want to ensure candidate p 's victory: It is enough to add all the voters who vote for p (or as many as we are allowed) and then, in a loop, keep deleting voters who vote for a candidate other than p with the highest score, until p is the only candidate with the highest score or we have exceeded our limit of voters to delete. Finally, in a loop, keep bribing voters who vote for a candidate other than p with the highest score to vote for p , until p is the only candidate with the highest score or we have exceeded our limit of voters to bribe. If p becomes a unique winner via this procedure, then accept. Otherwise reject. We omit the easy proof for the destructive case.

(b) Let I be an instance of destructive Condorcet-AC+AV+DV+BV, where our goal is to prevent candidate p from being a Condorcet winner (we assume that p is a Condorcet winner before any control action is performed). It is enough to ensure that some candidate c wins a head-to-head contest with p . Our algorithm works as follows.

Let C be the set of candidates originally in the election and let A be the set of candidates that we can add (we take $A = \emptyset$ if we are not allowed to add any candidates). For each $c \in (C \cup A) - \{p\}$ we do the following:

1. Add as many voters who prefer c to p as possible.
2. Delete as many voters who prefer p to c as possible.

⁸Constructive bribery for plurality and constructive bribery for approval have been considered in [FHH09a] and constructive and destructive bribery for Copeland has been studied in [FHHR09a]. In Theorem 3.7 we—in effect—give polynomial-time algorithms for destructive bribery in plurality, approval, and Condorcet. Constructive Condorcet-BV is NP-complete and this is implicitly shown in [FHHR09a, Theorem 3.2].

3. Among the remaining voters who prefer p to c , bribe as many as possible to rank c first.

If after these actions c wins his or her head-to-head contest with p then we accept. If no $c \in (C \cup A) - \{p\}$ leads to acceptance, then we reject. It is easy to see that this algorithm is correct and runs in polynomial time. (We point out that it is enough to add only a single candidate, the candidate c that prevents p from winning, if he or she happens to be a member of A).

For the case of approval, our algorithm works similarly, except the following differences: We add voters who approve of c but not of p . We delete voters who approve of p but not of c . For each remaining voter v_i , if we still have not exceeded our bribing limit, if v_i approves of p but not of c , we bribe v_i to reverse approvals on p and c . (Note that if we do not exceed our bribing limit by this procedure, this means that each voter that approves of p also approves of c and thus p is not a unique winner.) If these actions lead to p not being a unique winner, we accept. If we do not accept for any $c \in (C \cup A) - \{p\}$, we reject.

(c) The idea is to combine Copeland $^\alpha$ destructive-AC and destructive-DC algorithms [FHHR09a]. We give the full proof for the sake of completeness.

Let us fix a rational value α , $0 \leq \alpha \leq 1$. Given an election E and a candidate c in this election, we write $\text{score}_E^\alpha(c)$ to denote Copeland $^\alpha$ score of c . Let I be an instance of destructive Copeland $^\alpha$ -AC+DC control, with an election $E = (C, V)$, where we can add at most k_{AC} spoiler candidates from the set A , and where we can delete at most k_{DC} candidates. Our goal is to ensure that some despised candidate $d \in C$ is not a unique winner. Our algorithm is based on the following simple observation of Faliszewski et al. [FHHR09a]. For each candidate $c \in C$:

$$\text{score}_{(C,V)}^\alpha(c) = \sum_{c' \in C - \{c\}} \text{score}_{(\{c,c'\},V)}^\alpha(c).$$

Our goal is to prevent candidate d from being a unique winner. If d is not a unique winner, we immediately accept. Otherwise, we seek a candidate $c \in C \cup A$ such that we can ensure that c 's score is at least as high as that of d . Thus, for each $c \in C \cup A$ we do the following.

1. If $c \in A$, and $k_{\text{AC}} > 0$, we add c to the election (and if $c \in A$ but $k_{\text{AC}} = 0$, we proceed to the next c).
2. As long as we can still add more candidates, we keep executing the following operation: If there is a candidate $c' \in A$ such that value $a(c') = \text{score}_{(\{c,c'\},V)}^\alpha(c) - \text{score}_{(\{d,c'\},V)}^\alpha(d)$ is positive, we add a candidate $c'' \in A$, for whom $a(c'')$ is highest.
3. As long as we can still delete candidates, we keep executing the following operation: If there is a candidate $c' \in C$ such that value $r(c') = \text{score}_{(\{d,c'\},V)}^\alpha(d) - \text{score}_{(\{c,c'\},V)}^\alpha(c)$ is positive, we delete a candidate $c'' \in C$, for whom $r(c'')$ is highest.
4. If after these steps d is not a unique winner, we accept.

If we do not accept for any $c \in C \cup A$, we reject.

It is easy to see that we never delete a candidate that we have added. Also, it is easy to see that the algorithm works in polynomial time, and that it is correct. Correctness follows from the fact that (a) in the main loop of the algorithm, when dealing with candidate $c \in C \cup A$, each addition of a candidate and each deletion of a candidate increases the difference between the score of c and the score of d as much as is possible, and (b) the order of adding/deleting candidates is irrelevant. \square

As witnessed by Theorem 3.7 and the results of Section 4, for all natural election systems that we have considered, all constructive vulnerabilities combine and so do all destructive ones. It is natural to wonder whether this is a necessary consequence of our model of multiprong control or whether in fact there is an election system for which combining two control types to which the system is vulnerable yields a multipronged control problem to which the system is resistant. Theorem 3.8 shows that the latter is the case, even for a natural (though rather unusual) election system.

In the thirteenth century, Ramon Llull proposed an election system that could be used to choose popes and leaders of monastic orders (see [HP01, ML06]). In his system, voters choose the winner from among themselves (so, the candidates are the same as the voters). Apart from that, Llull's voting system is basically Copeland¹, the version of Copeland that most richly rewards ties. Formally, we define the voting system OriginalLlull as follows: For an election $E = (C, V)$, if the set of names of V , which we will denote by $\text{names}(V)$, is not equal to C , then there are no winners. Otherwise, a candidate $c \in C$ is a winner if and only if it is a Copeland¹ winner. Note that single-prong AC and AV control for OriginalLlull don't make all that much sense, and so it should come as no surprise that OriginalLlull is vulnerable to both constructive AC control and constructive AV control. In addition, we will show (by renaming and padding) that Copeland¹-AV can be reduced to OriginalLlull¹-AC+AV. Since Copeland¹ is resistant to constructive control by adding voters [FHR09a], this then leads to the following theorem.

Theorem 3.8. *OriginalLlull is vulnerable to both constructive AC control and constructive AV control but is resistant to constructive AC+AV control.*

Proof. It is immediate that OriginalLlull is susceptible to constructive AC, AV, and (by Theorem 3.3) AC+AV control. It is also easy to see that constructive OriginalLlull-AC (AV) control is in P: If possible add candidates (voters) such that the set of voter names is equal to the set of candidates, and then check if the preferred candidate is a unique Copeland¹ winner. If this is not possible, reject.

We will now show, via a reduction from constructive Copeland¹-AV control (which is NP-hard [FHR09a]) that constructive OriginalLlull-AC+AV control is NP-hard. Let C be a set of candidates, V and W be two disjoint collections of voters with preference lists over C , $p \in C$ the preferred candidate, and $k \in \mathbb{N}$. The question is whether there exists a subcollection $W' \subseteq W$ of size at most k such that p is a unique Copeland¹ winner of $(C, V \cup W')$. Without loss of generality, we assume that V is not empty.

We will now show how to pad this election. For an OriginalLlull election to be non-trivial, we certainly need to have the same number of candidates as voters (later, we will also rename the voters so that they are the same as the candidates). If $\|V\| < \|C\|$, we want to add a collection of new dummy voters V' such that $\|V\| + \|V'\| = \|C\|$ and such that adding V' to an election does not change the relative Copeland¹ scores of the candidates. This can be accomplished by letting half of the voters in V' vote C (recall Convention A) and half of the voters in V' vote \overleftarrow{C} . Of course, this can only be done if $\|V'\|$ is even.

So, we will do the following. If $\|V\| < \|C\|$, we add a collection of new voters V' such that $\|V'\| = \|C\| - \|V\|$ if $\|C\| - \|V\|$ is even, and $\|V'\| = \|C\| - \|V\| + 1$ if $\|C\| - \|V\|$ is odd. If $\|V\| \geq \|C\|$, we let $V' = \emptyset$. Half of the voters in V' vote C and half of the voters in V' vote \overleftarrow{C} . In addition, we introduce a set A of new candidates such that $\|C\| + \|A\| = \|V\| + \|V'\| + \|W\|$. Note that this is always possible, since $\|V\| + \|V'\| \geq \|C\|$. We extend the votes of the voters (in V , V' , and W) to $C \cup A$ by taking their preference order on C and following this by the candidates in A in some fixed, arbitrary order. Note that this will have the effect that candidates in A will never be winners.

Let $W' \subseteq W$, $A' \subseteq A$, $E = (C, V \cup W')$, $E' = (C \cup A', V \cup V' \cup W')$. It is easy to see that the following hold (recall that V is not empty).

1. For all $d \in A'$, $\text{score}_{E'}^1(d) \leq \|A'\| - 1$.
2. For all $c \in C$, $\text{score}_{E'}^1(c) = \text{score}_E^1(c) + \|A'\|$.
3. For all $c, c' \in C, c \neq c'$, $\text{score}_E^1(c) - \text{score}_E^1(c') = \text{score}_{E'}^1(c) - \text{score}_{E'}^1(c')$.
4. p is a unique Copeland¹ winner of E if and only if p is a unique Copeland¹ winner of E' .

We are now ready to define the reduction. Name the voters such that $\text{names}(V \cup V') \supseteq C$ and $\text{names}(V \cup V' \cup W) = C \cup A$. Then map (C, V, W, p, k) to $(C, A, V \cup V', W, p, \|A\|, k)$. We claim that p can be made a unique Copeland¹ winner of (C, V) by adding at most k voters from W if and only if p can be made a unique OriginalLlull winner of $(C, V \cup V')$ by adding (an unlimited number of) candidates from A and at most k voters from W .

First suppose that W' is a subcollection of W of size at most k such that p is the unique Copeland¹ winner of $(C, V \cup W')$. Let $A' \subseteq A$ be the set of candidates such that $C \cup A' = \text{names}(V \cup V' \cup W')$. By item 4 above, p is the unique Copeland¹ winner of $(C \cup A', V \cup V' \cup W')$, and thus p is the unique OriginalLlull winner of $(C \cup A', V \cup V' \cup W')$.

For the converse, suppose that there exist $A' \subseteq A$ and $W' \subseteq W$ such that $\|W'\| \leq k$, and p is the unique OriginalLlull winner of $(C \cup A', V \cup V' \cup W')$. Then p is the unique Copeland¹ winner of $(C \cup A', V \cup V' \cup W')$, and, by item 4, p is the unique Copeland¹ winner of $(C, V \cup W')$.

Thus our reduction is correct and, since it can be computed in polynomial time, the proof is complete. \square

OriginalLlull is neutral (permuting the names of the candidates does not affect the outcome of an election) but not anonymous (renaming the voters can change the outcome

of an election). By sneakily building the preference orders of the voters into the names of the candidates, we can make the system anonymous as well as neutral (at the price of losing naturalness).

Theorem 3.9. *There exists a neutral and anonymous election system \mathcal{E} such that \mathcal{E} is vulnerable to both constructive AC control and constructive AV control but is resistant to constructive AC+AV control.*

Proof. We first describe \mathcal{E} . On input (C, V) , an election, if there exists a set $I \subseteq \mathbb{N}^+$ and bijections c from I to C and v from I to V such that for all $i \in I$, $c(i) = (i, >_i)$ where $>_i$ is a preference order on I (i.e., we interpret candidate names as pairs consisting of a positive integer and a preference order on I) and voter $v(i)$ corresponds to candidate $c(i)$ in the sense that for all $j, k \in I$, $j >_i k$ if and only if $c(j) > c(k)$ in voter $v(i)$'s preference order, then the winners are exactly the Copeland¹ winners. Otherwise, there are no winners.

Note that \mathcal{E} is neutral and anonymous and basically the same as OriginalLlull. The same argument as used for OriginalLlull in the proof of Theorem 3.8 shows that \mathcal{E} is vulnerable to constructive AC and AV control and susceptible to AC+AV control. To show that constructive \mathcal{E} -AC+AV control is NP-hard, we adapt the reduction from constructive Copeland¹-AV control to constructive OriginalLlull-AC+AV control from the proof of Theorem 3.8. Let C be a set of candidates, V and W be two disjoint collections of voters with preference lists over C , $p \in C$ the preferred candidate, and $k \in \mathbb{N}$. Without loss of generality, we assume that V is not empty. Let V' and A be as in the proof of Theorem 3.8. Recall that $\|V \cup V'\| \geq \|C\|$ and $\|V \cup V' \cup W\| = \|C \cup A\|$. From the proof of Theorem 3.8 we have the following.

Claim 3.10. *Let $W' \subseteq W$, $A' \subseteq A$, $E = (C, V \cup W')$, $E' = (C \cup A', V \cup V' \cup W')$. p is a unique Copeland¹ winner of E if and only if p is a unique Copeland¹ winner of E' .*

We are now ready to define the reduction. We will first rename the candidates. Note that renaming candidates does not change the outcome of a Copeland¹ election. Number the candidates in $C \cup A$ from 1 to $\|C \cup A\|$ such that the candidates in C are numbered from 1 to $\|C\|$. Number the voters in $V \cup V' \cup W$ from 1 to $\|V \cup V' \cup W\|$ ($= \|C \cup A\|$) such that the voters in $\|V \cup V'\|$ are numbered from 1 to $\|V \cup V'\|$. Now rename candidate i to $c_i = (i, >_i)$ where $>_i$ is the preference order on $\{1, \dots, \|C\| + \|A\|\}$ such that for all $j, k \in \{1, \dots, \|C\| + \|A\|\}$, $j >_i k$ if and only if $j > k$ in voter i . Rename all candidates occurring in C , A , V , V' , and W in this way. We claim that p can be made a unique Copeland¹ winner of (C, V) by adding at most k voters from W if and only if p can be made a unique \mathcal{E} winner of $(C, V \cup V')$ by adding candidates from A and at most k voters from W .

First suppose that W' is a subcollection of W of size at most k such that p is the unique Copeland¹ winner of $(C, V \cup W')$. Let $A' \subseteq A$ be the set of candidates such that $C \cup A' = \{c_i \mid \text{voter } i \text{ is in } V \cup V' \cup W'\}$. By Claim 3.10, p is the unique Copeland¹ winner of $(C \cup A', V \cup V' \cup W')$, and thus p is the unique \mathcal{E} winner of $(C \cup A', V \cup V' \cup W')$.

For the converse, suppose that there exist $A' \subseteq A$ and $W' \subseteq W$ such that $\|W'\| \leq k$, and p is the unique \mathcal{E} winner of $(C \cup A', V \cup V' \cup W')$. Then p is the unique Copeland¹

winner of $(C \cup A', V \cup V' \cup W')$, and by Claim 3.10, p is the unique Copeland¹ winner of $(C, V \cup W')$.

Thus our reduction is correct and, since it can be computed in polynomial time, the proof is complete. \square

4 Control in Maximin

In this section we initiate the study of control in the maximin election system. Maximin is loosely related to Copeland ^{α} voting in the sense that both are defined in terms of the pairwise head-to-head contests. In addition, the unweighted coalitional manipulation problem for maximin and Copeland ^{α} ($\alpha \neq 0.5$) exhibits the same unusual behavior: It is in P for one manipulator and NP-complete for two or more manipulators [XCP09, FHS08, FHS10]. Thus one might wonder whether both systems will be similar with regard to their resistances to control. In fact, there are very interesting differences.

It is easy to see that maximin is susceptible to all basic types of constructive and destructive control. And so, by Theorem 3.3, to show vulnerability to constructive (destructive) \mathcal{C} control it suffices to give a polynomial-time algorithm that decides the constructive (destructive) $\mathcal{E}\text{-}\mathcal{C}$ problem, and to show resistance to constructive (destructive) \mathcal{C} control it suffices to show that the constructive (destructive) $\mathcal{E}\text{-}\mathcal{C}$ problem is NP-hard.

4.1 Candidate Control in Maximin

Let us now focus on candidate control in maximin, that is, on AC, AC_u , and DC control types, both in the constructive and in the destructive setting. As is the case for Copeland ^{α} , $0 \leq \alpha \leq 1$, maximin is resistant to control by adding candidates.

Theorem 4.1. *Maximin is resistant to constructive AC control.*

Proof. We give a reduction from X3C. Let (B, \mathcal{S}) , where $B = \{b_1, \dots, b_{3k}\}$ is a set of $3k$ elements and $\mathcal{S} = \{S_1, \dots, S_n\}$ is a set of 3-subsets of B , be our input X3C instance. We form an election $E = (C \cup A, V)$, where $C = B \cup \{p\}$, $A = \{a_1, \dots, a_n\}$, and $V = (v_1, \dots, v_{2n+2})$. (Candidates in A are the spoiler candidates, which the attacker has the ability to add to election (C, V) .)

Voters in V have the following preferences. For each $S_i \in \mathcal{S}$, voter v_i reports preference list $p > \overbrace{B - S_i} > a_i > S_i > A - \{a_i\}$ and voter v_{n+i} reports preference list $\overbrace{A - \{a_i\}} > a_i > \overbrace{S_i} > \overbrace{B - S_i} > p$. Voter v_{2n+1} reports $p > A > B$ and voter v_{2n+2} reports $\overbrace{B} > p > \overbrace{A}$.

We claim that there is a set $A' \subseteq A$ such that $\|A'\| \leq k$ and p is a unique winner of $(C \cup A', V)$ if and only if (B, \mathcal{S}) is a “yes”-instance of X3C.

To show the claim, let $E' = (C, V)$. For each pair of distinct elements $b_i, b_j \in B$, we have that $N_{E'}(b_i, b_j) = n + 1$, $N_{E'}(p, b_i) = n + 1$, and $N_{E'}(b_i, p) = n + 1$. That is, all candidates in E' tie. Now consider some set $A'' \subseteq A$, $\|A''\| \leq k$, and an election $E'' = (C \cup A'', V)$. Values of $N_{E''}$ and $N_{E'}$ are the same for each pair of candidates in $\{p\} \cup B$. For each pair of distinct

elements $a_i, a_j \in A''$, we have $N_{E''}(p, a_i) = n + 2$, $N_{E''}(a_i, p) = n$, and $N_{E''}(a_i, a_j) = n + 1$. For each $b_i \in B$ and each $a_j \in A''$ we have that

$$N_{E''}(b_i, a_j) = \begin{cases} n & \text{if } b_i \in S_j, \\ n + 1 & \text{if } b_i \notin S_j, \end{cases}$$

and, of course, $N_{E''}(a_j, b_i) = 2n + 2 - N_{E''}(b_i, a_j)$. Thus, by definition of maximin, we have the following scores in E'' : (a) $\text{score}_{E''}(p) = n + 1$, (b) for each $a_j \in A''$, $\text{score}_{E''}(a_j) = n$, and (c) for each $b_i \in B$,

$$\text{score}_{E''}(b_i) = \begin{cases} n & \text{if } (\exists a_j \in A'') [b_i \in S_j], \\ n + 1 & \text{otherwise.} \end{cases}$$

A'' corresponds to a family S'' of 3-sets from \mathcal{S} such that for each j , $1 \leq j \leq n$, S'' contains set S_j if and only if A'' contains a_j . Since $\|A''\| \leq k$, it is easy to see that p is a unique winner of E'' if and only if S'' is an exact cover of B . \square

Copeland $^\alpha$, $0 \leq \alpha \leq 1$, is resistant to constructive AC control, but for $\alpha \in \{0, 1\}$, Copeland $^\alpha$ is vulnerable to constructive control by adding an unlimited number of candidates. It turns out that so is maximin. However, interestingly, in contrast to Copeland, maximin is also vulnerable to DC control, and in fact even to AC_u+DC control. Intuitively, in constructive AC_u+DC control we should add as many candidates as possible (because adding a candidate generally decreases other candidates' scores, making our preferred candidate's way to victory easier) and then delete those candidates who stand in our candidate's way (i.e., those whose existence blocks the preferred candidate's score from increasing). Studying constructive AC_u+DC control for maximin jointly leads to a compact, coherent algorithm. If we were to consider both control types separately, we would have to give two fairly similar algorithms while obtaining a weaker result.

Theorem 4.2. *Maximin is vulnerable to constructive AC_u+DC control.*

Proof. We give a polynomial-time algorithm for constructive maximin-AC_u+DC control. The input contains an election $E = (C, V)$, a set of spoiler candidates A , a preferred candidate $p \in C$, and a nonnegative integer k_{DC} . Voters in V have preference lists over the candidates in $C \cup A$. We ask whether there exist sets $A' \subseteq A$ and $C' \subseteq C$ such that (a) $\|C'\| \leq k_{DC}$ and (b) p is a unique winner of election $((C - C') \cup A', V)$. If $k_{DC} \geq \|C\| - 1$, we accept immediately because we can delete all candidates but p . Otherwise, we use the following algorithm.

Preparation. We rename the candidates in C and A so that $C = \{p, c_1, \dots, c_m\}$ and $A = \{c_{m+1}, \dots, c_{m+m'}\}$. Let $E' = (C \cup A, V)$ and let $P = \{N_{E'}(p, c_i) \mid c_i \in C \cup A\}$. That is, P contains all the values that candidate p may obtain as scores upon deleting some candidates from E' . For each $k \in P$, let $Q(k) = \{c_i \mid c_i \in C \cup A - \{p\} \wedge N_{E'}(p, c_i) < k\}$. Intuitively, $Q(k)$ is the set of candidates in E' that prevent p from having at least k points.

Main loop. For each $k \in P$, our algorithm tests whether by deleting at most k_{DC} candidates from C and any number of candidates from A it is possible to ensure that p obtains exactly k points and becomes a unique winner of E' . Let us fix some value $k \in P$. We build a set D of candidates to delete. Initially, we set $D = Q(k)$. It is easy to see that deleting candidates in $Q(k)$ is a necessary and sufficient condition for p to have score k . However, deleting candidates in $Q(k)$ is not necessarily sufficient to ensure that p is a unique winner because candidates with scores greater or equal to k may exist. We execute the following loop (which we will call the *fixing loop*):

1. Set $E'' = ((C \cup A) - D, V)$.
2. Pick a candidate $d \in (C \cup A) - D$ such that $\text{score}_{E''}(d) \geq k$ (break from the loop if no such candidate exists).
3. Add d to D and jump back to Step 1.

We accept if $C \cap D \leq k_{DC}$ and we proceed to the next value of k otherwise.⁹ If none of the values $k \in P$ leads to acceptance then we reject.

Let us now briefly explain why the above algorithm is correct. It is easy to see that in maximin adding some candidate c to an election does not increase other candidates' scores, and deleting some candidate d from an election does not decrease other candidates' scores. Thus, if after deleting candidates in $Q(k)$ there still are candidates other than p with k points or more, the only way to ensure p 's victory—without explicitly trying to increase p 's score—is by deleting those candidates. Also, clearly, the only way to ensure that p has exactly k points is by deleting candidates $Q(k)$.

Note that during the execution of the fixing loop, the score of p might increase to some value $k' > k$. If that happens, it means that it is impossible to ensure p 's victory while keeping his or her score equal to k . However, we do not need to change k to k' in that iteration of the main loop as we will consider k' in a different iteration. \square

Maximin is also vulnerable to destructive AC+DC control. The proof relies on the fact that (a) if there is a way to prevent a despised candidate from winning a maximin election via adding some spoiler candidates then there is a way to do so by adding at most 2 candidates, (b) adding a candidate cannot increase the score of any candidate other than the added one, and (c) deleting a candidate cannot decrease the score of any candidate other than the deleted one. In essence, the algorithm performs a brute-force search for the candidates to add and then uses the constructive maximin-DC control algorithm from Theorem 4.2.

Theorem 4.3. *Maximin is vulnerable to destructive AC+DC control.*

⁹If we accept, D implicitly describes the control action that ensures p 's victory: We should delete from C the candidates in $C \cap D$ and add from A the candidates in $A - D$.

Proof. We will first give an algorithm for destructive maximin-AC and then argue how it can be combined with the algorithm from Theorem 4.2 to solve destructive maximin-AC+DC in polynomial time.

Let us first focus on the destructive AC problem. Our input is an election $E = (C, V)$, where $C = \{d, c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$, a spoiler candidate set $A = \{c_{m+1}, \dots, c_{m'}\}$, and a nonnegative integer k_{AC} . The voters have preference orders over $C \cup A$. The goal is to ensure that d is not a unique winner of E via adding at most k_{AC} candidates from A .

Let us assume that there exists a set $A' \subseteq A$ such that d is not a unique winner of election $E' = (C \cup A', V)$. Since d is not a unique winner of E' , there exists some candidate $c' \in C \cup A'$ such that $\text{score}_{E'}(c') \geq \text{score}_{E'}(d)$. Also, by definition of maximin, there is some candidate $d' \in C \cup A'$ such that $\text{score}_{E'}(d) = N_{E'}(d, d')$. As a consequence, d is not a unique winner of election $E'' = (C \cup \{c', d'\}, V)$. The reason is that $\text{score}_{E''}(d) = \text{score}_{E'}(d)$ (because both E' and E'' contain d') and $\text{score}_{E''}(c') \geq \text{score}_{E'}(c')$ (because adding the remaining $A' - \{c', d'\}$ candidates to E'' does not increase c' 's score). Thus, to test whether it is possible to ensure that d is not a unique winner of E , it suffices to test whether there is a set $A'' \subseteq A$ such that $\|A''\| \leq \min(2, k_{AC})$ and d is not a unique winner of $(C \cup A'', V)$. Clearly, this test can be carried out in polynomial time.

Let us now consider the AC+DC case. The input and the goal are the same as before, except that now we are also given a nonnegative integer k_{DC} and we are allowed to delete up to k_{DC} candidates. We now describe our algorithm. For each set $\{c', d'\}$ of up to 2 candidates, $\{c', d'\} \subseteq (C \cup A) - \{d\}$ we execute the following steps.

1. We check if $\|A \cap \{c', d'\}\| \leq k_{AC}$ (and we proceed to the next $\{c', d'\}$ if this is not the case).
2. We compute a set $D \subseteq C - \{d, c', d'\}$, $\|D\| \leq k_{DC}$, that maximizes $\text{score}_{E'}(c')$, where $E' = ((C \cup \{c', d'\}) - D, V)$.
3. If d is not a unique winner of $E' = ((C \cup \{c', d'\}) - D, V)$, we accept.

We reject if we do not accept for any $\{c', d'\} \subseteq (C \cup A) - \{d\}$.

The intended role of d' is to lower the score of d and keep it at a fixed level, while, of course, the intended role of c' is to defeat d . By reasoning analogous to that for the AC case, we can see that there is no need to add more than two candidates. Thus, given $\{c', d'\}$, it remains to compute the appropriate set D . In essence, we can do so in the same manner as in the constructive AC+DC case.

Let k be some positive integer. We set $D(k) = \{c_i \in C - \{c', d', d\} \mid N_E(c', c_i) < k\}$ and we pick $D = D(i)$, where i is as large as possible (but no larger than $\|V\|\)$ and $\|D\| \leq k_{DC}$. Deleting candidates in D maximizes the score of c' , given that we cannot delete d and d' . It is easy to see that this D can be computed in polynomial time. \square

4.2 Control by Adding and Deleting Voters in Maximin

In this section we consider the complexity of constructive and destructive AV and DV control types. (We will consider bribery, BV, in the next section; recall that in this paper, bribery is a basic control type, though it is usually treated separately in the literature.) In the previous section we have seen that maximin is vulnerable to all basic types of constructive and destructive candidate control except for constructive control by adding candidates (constructive AC control). The situation regarding voter control is quite different: As shown in the next three theorems, maximin is resistant to all basic types of constructive and destructive voter control.

Theorem 4.4. *Maximin is resistant to constructive and destructive AV control.*

Proof. We will first give an NP-hardness proof for the constructive case and then we will describe how to modify it for the destructive case.

We now give a reduction of the X3C problem to the constructive maximin-AV problem. Our input X3C instance is (B, \mathcal{S}) , where $B = \{b_1, \dots, b_{3k}\}$ is a set of $3k$ distinct elements and $\mathcal{S} = \{S_1, \dots, S_n\}$ is a family of n 3-element subsets of B . Without loss of generality, we assume $k \geq 1$. Our reduction outputs the following instance. We have an election $E = (C, V)$, where $C = B \cup \{p, d\}$ and $V = (v_1, \dots, v_{4k})$. There are $2k$ voters with preference order $d > B > p$, k voters with preference order $p > B > d$, and k voters with preference order $p > d > B$. In addition, we have a collection $W = (w_1, \dots, w_n)$ of unregistered voters, where the i 'th voter, $1 \leq i \leq n$, has preference order

$$B - S_i > p > S_i > d.$$

We claim that there is a subcollection $W' \subseteq W$ such that $\|W'\| \leq k$ and p is a unique winner of election $(C, V \cup W')$ if and only if (B, \mathcal{S}) is a "yes"-instance of X3C.

It is easy to verify that for each $b_i \in B$ it holds that $N_E(p, b_i) = 2k$, and that $N_E(p, d) = 2k$. Thus, $\text{score}_E(p) = 2k$. Similarly, it is easy to verify that $\text{score}_E(d) = 2k$, and that for each $b_i \in B$, $\text{score}_E(b_i) \leq k$. Let W'' be a subcollection of W such that $\|W''\| \leq k$ and let $E'' = (C, V \cup W'')$. For each $b_i \in B$ it holds that $\text{score}_{E''}(b_i) \leq 2k$. Since each voter in W ranks d as the least desirable candidate, $\text{score}_{E''}(d) = 2k$. What is p 's score in election E'' ? If there exists a candidate $b_i \in B$ such that there is no voter w_j in W'' that prefers p to b_i , then $\text{score}_{E''}(p) = 2k$ (because $N_{E''}(p, b_i) = 2k$). Otherwise, $\text{score}_{E''}(p) \geq 2k + 1$. Thus, p is a unique winner of E'' if and only if W'' corresponds to an exact cover of B . This proves our claim and, as the reduction is clearly computable in polynomial time, concludes the proof for the constructive maximin-AC case.

To show that destructive maximin-AC is NP-hard, we use the same reduction, except that we remove from V a single voter with preference list $p > B > d$, and we set the task to preventing d from being a unique winner. Removing a $p > B > d$ voter from V ensures that before we start adding candidates, d has score $2k$ (and this score cannot be changed), p has score $2k - 1$ (and p needs to get one point extra over each other candidate to increase his or her score and prevent d from being a unique winner), and each $b_i \in B$ has score

$k - 1$ (thus, no candidate in B can obtain score higher than $2k - 1$ via adding no more than k candidates from W). The same reasoning as for the constructive case proves that the reduction correctly reduces X3C to destructive maximin-AV. \square

Theorem 4.5. *Maximin is resistant to constructive and destructive DV control.*

Proof. We will first show NP-hardness for constructive maximin-DV control and then we will argue how to modify the construction to obtain the result for the destructive case.

Our reduction is from X3C. Let (B, \mathcal{S}) be our input X3C instance, where $B = \{b_1, \dots, b_{3k}\}$, $\mathcal{S} = \{S_1, \dots, S_n\}$, and for each i , $1 \leq i \leq n$, $\|S_i\| = 3$. Without loss of generality, we assume that $n \geq k \geq 3$ (if $n < k$ then \mathcal{S} does not contain a cover of B , and if $k \leq 2$ we can solve the problem by brute force). We form an election $E = (C, V)$, where $C = B \cup \{p, d\}$ and where $V = V' \cup V''$, $V' = (v'_1, \dots, v'_{2n})$, $V'' = (v''_1, \dots, v''_{2n-k+2})$. For each i , $1 \leq i \leq n$, voter v'_i has preference order

$$d > B - S_i > p > S_i$$

and voter v'_{n+i} has preference order

$$d > \overleftarrow{S_i} > p > \overleftarrow{B - S_i}.$$

Among the voters in V'' we have: 2 voters with preference order $p > d > B$, $n - k$ voters with preference order $p > B > d$, and n voters with preference order $B > p > d$. We claim that it is possible to ensure that p is a unique winner of election E via deleting at most k voters if and only if (B, \mathcal{S}) is a “yes”-instance of X3C.

Via routine calculation we see that candidates in election E have the following scores:

1. $\text{score}_E(d) = 2n$ (because $N_E(d, p) = 2n$ and for each $b_i \in B$, $N_E(d, b_i) = 2n + 2$),
2. $\text{score}_E(p) = 2n - k + 2$ (because $N_E(p, d) = 2n - k + 2$ and for each $b_i \in B$, $N_E(p, b_i) = 2n - k + 2$), and
3. for each $b_i \in B$, $\text{score}_E(b_i) \leq 2n - k$ (because $N_E(b_i, d) = 2n - k$).

Before any voters are deleted, d is the unique winner with $k - 2$ more points than p . Via deleting at most k voters it is possible to decrease d 's score at most by k points. Let W be a collection of voters such that p is the unique winner of $E' = (C, V - W)$. We partition W into $W' \cup W''$, where W' contains those members of W that belong to V' and W'' contains those members of W that belong to V'' . We claim that W'' is empty. For the sake of contradiction let us assume that $W'' \neq \emptyset$. Let $E'' = (C, V - W'')$. Since every voter in V'' prefers p to d , we have that $N_{E''}(p, d) = N_E(p, d) - \|W''\|$ and, as a result, $\text{score}_{E''}(p) \leq \text{score}_E(p) - \|W''\|$. In addition, assuming W'' is not empty, it is easy to observe that $\text{score}_{E''}(d) \geq \text{score}_E(d) - \|W''\| + 1$ (the reason for this is that deleting any single member of V'' does not decrease d 's score). That is, we have that:

$$\begin{aligned} \text{score}_{E''}(p) &\leq 2n - k + 2 - \|W''\|, \\ \text{score}_{E''}(d) &\geq 2n + 1 - \|W''\|. \end{aligned}$$

So in E'' , d has at least $k - 1$ more points than p . Since $\|W''\| \geq 1$, we can delete at most $k - 1$ voters W' from election E'' . But then p will not be a unique winner of E' , which is a contradiction.

Thus, W contains members of V' only. Since d is ranked first in every vote in V' , deleting voters from W decreases d 's score by exactly $\|W\|$. Further, deleting voters W certainly decreases p 's score by at least one point. Thus, after deleting voters W we have:

1. $\text{score}_{E'}(d) = 2n - \|W\|$,
2. $\text{score}_{E'}(p) \leq 2n - k + 2 - 1 = 2n - k + 1$.

In consequence, the only possibility that p is a unique winner after deleting voters W is that $\|W\| = k$ and we have equality in item 2 above. It is easy to verify that this equality holds if and only if W contains k voters among v'_1, \dots, v'_n that correspond to an exact cover of B via sets from \mathcal{S} (recall that $k \geq 3$). This proves that our reduction is correct, and since the reduction is clearly computable in polynomial time, completes the proof of NP-hardness of constructive maximin-DV control.

Let us now consider the destructive case. Let (B, \mathcal{S}) be our input X3C instance (with B and \mathcal{S} as in the constructive case). We form election $E = (C, V)$ which is identical to the one created in the constructive case, except that $V'' = (v''_1, \dots, v''_{2n-k})$ and we set these voters' preference orders as follows: There is one voter with preference order $p > d > B$, $n - k$ voters with preference order $p > B > d$, and $n - 1$ voters with preference order $B > p > d$. (That is, compared to the constructive case, we remove one voter with preference order $p > d > B$ and one with preference order $B > p > d$.) It is easy to see that d is the unique winner of election E and we claim that he or she can be prevented from being a unique winner via deleting at most k voters if and only if there is an exact cover of B by k sets from \mathcal{S} .

Via routine calculation, it is easy to verify that $\text{score}_E(d) = 2n$, and that $\text{score}_E(p) = 2n - k$. The former holds because $N_E(d, p) = 2n$ and $N_E(d, b_i) = 2n + 1$ and the latter holds because $N_E(p, d) = 2n - k$ and for each candidate $b_i \in B$ we have $N_E(p, b_i) = 2n - k + 1$. In addition, each candidate $b_i \in B$ has score at most $2n - k - 1$. Thus, it is possible to ensure that d is not a unique winner via deleting at most k voters if and only if there are exactly k voters deleting whom would decrease the score of d by k points and would not decrease p 's score. Let us assume that such a collection of voters exists and let W be such a collection. Since every voter in V'' prefers p to d , clearly W does not contain any voter in V'' . Thus, W contains exactly k voters from V' . Since for each $b_i \in B$ we have $N_E(p, b_i) = 2n - k + 1$, for each $b_i \in B$ W contains at most one voter who prefers p to b_i . Since $\|B\| = 3k$ and $k \geq 3$, this implies that W contains exactly a collection of voters corresponding to some exact cover of B by sets in \mathcal{S} . This completes the proof for the destructive case. \square

4.3 Bribery in Maximin

We now move on to bribery in maximin. Given the previous results, it is not surprising that maximin is resistant both to constructive bribery and to destructive bribery. Our proof is

an application of the “UV technique” of Faliszewski et al. [FHHR09a]. Very informally, the idea is to build an election in a way that ensures that the briber is limited to bribing only those voters who rank two special candidates ahead of the preferred one.

Theorem 4.6. *Maximin is resistant to constructive and destructive BV control.*

Proof. Our proofs follow via reductions from X3C. The reduction for the constructive case is almost identical the one for the constructive case and thus we will consider both cases in parallel.

Our reductions work as follows. Let (B, \mathcal{S}) be an instance of X3C, where $B = \{b_1, \dots, b_{3k}\}$ is a set of $3k$ distinct elements, and $\mathcal{S} = \{S_1, \dots, S_n\}$ is a family of 3-element subsets of B . (Without loss of generality, we assume that $n > k > 1$. If this is not the case, it is trivial to verify if (B, \mathcal{S}) is a “yes” instance of X3C.) We construct a set of candidates $C = \{p, d, s\} \cup B$, where p is our preferred candidate (the goal in the constructive setting is to ensure p is a unique winner) and d is our despised candidate (the goal in the destructive setting is to prevent d from being a unique winner). We construct six collections of voters, $V^1, V^2, V^3, V^4, V^5, V^6$, as follows:

1. V^1 contains $2n$ voters, v_1^1, \dots, v_{2n}^1 . For each i , $1 \leq i \leq n$, voters v_i^1 and v_{i+n}^1 have the following preference orders:

$$\begin{aligned} v_i^1 &: d > s > S_i > p > B - S_i \\ v_{n+i}^1 &: \overleftarrow{B - S_i} > p > \overleftarrow{S_i} > d > s. \end{aligned}$$

2. V^2 contains $2k$ voters, v_1^2, \dots, v_{2k}^2 . For each i , $1 \leq i \leq k$, voters v_i^2 and v_{i+k}^2 have the following preference orders:

$$\begin{aligned} v_i^2 &: s > d > p > B \\ v_{k+i}^2 &: \overleftarrow{B} > d > p > s. \end{aligned}$$

3. V^3 contains $2k$ voters, v_1^3, \dots, v_{2k}^3 . For each i , $1 \leq i \leq k$, voters v_i^3 and v_{i+k}^3 have the following preference orders:

$$\begin{aligned} v_i^3 &: d > s > p > B \\ v_{k+i}^3 &: \overleftarrow{B} > s > p > d. \end{aligned}$$

4. V^4 contains $4k$ voters, v_1^4, \dots, v_{4k}^4 . For each i , $1 \leq i \leq 2k$, voters v_i^4 and v_{i+2k}^4 have the following preference orders:

$$\begin{aligned} v_i^4 &: d > B > p > s \\ v_{2k+i}^4 &: s > p > d > \overleftarrow{B}. \end{aligned}$$

| (a) Values of $N_{E_c}(\cdot, \cdot)$. | | | | |
|---|--------------|--------------|---------------|-------------------|
| | p | d | s | B |
| p | – | $n + 3k + 2$ | $n + 3k + 2$ | $n + 4k + 1$ |
| d | $n + 5k + 1$ | – | $2n + 4k + 2$ | $n + 6k + 2$ |
| s | $n + 5k + 1$ | $4k + 1$ | – | $n + 4k + 2$ |
| B | $n + 4k + 2$ | $n + 2k + 1$ | $n + 4k + 1$ | $\leq n + 4k + 2$ |

| (b) Values of $N_{E_d}(\cdot, \cdot)$. | | | | |
|---|--------------|--------------|---------------|--------------|
| | p | d | s | B |
| p | – | $n + 3k + 1$ | $n + 3k + 1$ | $n + 4k$ |
| d | $n + 5k + 1$ | – | $2n + 4k + 1$ | $n + 6k + 1$ |
| s | $n + 5k + 1$ | $4k + 1$ | – | $n + 4k + 1$ |
| B | $n + 4k + 2$ | $n + 2k + 1$ | $n + 4k + 1$ | $n + 4k + 1$ |

Table 1: Values of $N_{E_c}(\cdot, \cdot)$ and $N_{E_d}(\cdot, \cdot)$ for each pair of candidates. Let E be one of E_c, E_d . An entry in row $c' \in \{p, d, s\}$ and column $c'' \in \{p, d, s\}$, $c' \neq c''$, of the appropriate table above gives value $N_E(c', c'')$. For row B and for column B we adopt the following convention. For each $c \in \{p, d, s\}$ and for each $b_i \in B$, an entry in row B and column c is equal to $N_E(b_i, c)$. For each $c \in \{p, d, s\}$ and for each $b_i \in B$, an entry in row c and column B is equal to $N_E(c, b_i)$. For each two distinct $b_i, b_j \in B$, the entry in row B and column B is the upper bound on $N_E(b_i, b_j)$. (For E_d this entry is, in fact, exact.)

5. V^5 contains 2 voters, v_1^5, v_2^5 with the following preference orders

$$\begin{aligned} v_1^5 &: s > B > p > d \\ v_2^5 &: d > \overleftarrow{B} > p > s. \end{aligned}$$

6. V^6 contains a single voter, v_1^6 , with preference order $p > d > s > B$.

We form two elections, E_c and E_d , where $E_c = (C, V^1 \cup \dots \cup V^6)$ and $E_d = (C, V^1 \cup \dots \cup V^5)$; that is, E_c and E_d are identical except E_d does not contain the single voter from V^6 . E_c contains $2n + 8k + 3$ voters and E_d contains $2n + 8k + 2$ voters. Values of N_{E_c} and N_{E_d} for each pair of candidates are given in Table 1.

For the constructive case, we claim that it is possible to ensure that p is a unique winner of election E_c by bribing at most k voters if and only if (B, \mathcal{S}) is a “yes” instance of X3C. Let us now prove this claim. By inspecting Table 1, and recalling that $n > k > 1$, we see that $\text{score}_{E_c}(p) = n + 3k + 2$, $\text{score}_{E_c}(d) = n + 5k + 1$, $\text{score}_{E_c}(s) = 4k + 1$, and for each $b_i \in B$, $\text{score}_{E_c}(b_i) \leq n + 2k + 1$. That is, prior to any bribing, d is the unique winner and p has the second highest score.

It is easy to see that by bribing $t \leq k$ voters, the briber can change each candidate’s score by at most t points. Thus, for the bribery to be successful, the briber has to bribe exactly k voters in such a way that d ’s score decreases to $n + 4k + 1$ and p ’s score increases

to $n + 4k + 2$. To achieve this, the briber has to find a collection V' of voters such that $\|V'\| = k$, and

1. each voter in V' ranks p below both d and s , and
2. for each $b_i \in B$, there is a voter in V' who ranks p below b_i .

The only voters that satisfy the first condition are $v_1^1, \dots, v_n^1, v_1^2, \dots, v_k^2, v_1^3, \dots, v_k^3$. Further, among these voters only v_1^1, \dots, v_n^1 rank p below some member of B and, in fact, for each i , $1 \leq i \leq n$, v_i^1 ranks p below exactly three members of B . Thus, it is easy to see that each k voters from $v_1^1, \dots, v_n^1, v_1^2, \dots, v_k^2, v_1^3, \dots, v_k^3$ that satisfy the second condition correspond naturally to a cover of B by sets from \mathcal{S} . (Note that it suffices that the briber bribes voters in V' to rank p first without changing the votes in any other way, and that changing the votes in any other way than ranking p first is not necessary.) As a result, if it is possible to ensure that p is a winner of E_c by bribing at most k voters then (B, \mathcal{S}) is a “yes” instance of X3C. For the other direction, it is easy to verify that if (B, \mathcal{S}) is a “yes” instance of X3C then bribing k voters from v_1^1, \dots, v_n^1 that correspond to a cover of B to rank p first suffices to ensure that p is a unique winner. This completes the proof for the constructive case.

For the destructive case, we claim that it is possible to ensure that d is not a unique winner of E_d if and only if (B, \mathcal{S}) is a “yes” instance of X3C. The proof is analogous to the constructive case: It suffices to note that p is the only candidate that can possibly tie for victory with d . The rest of the proof proceeds as for the constructive case. \square

4.4 Connection to Dodgson Voting

We conclude our discussion of (control in) maximin voting with a small detour, showing a connection between maximin and the famous voting rule (i.e., election system) of Dodgson.

Dodgson voting, proposed in the 19th century by Charles Lutwidge Dodgson,¹⁰ works as follows [Dod76]. Let $E = (C, V)$ be an election, where $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$. For a candidate $c_i \in C$, the Dodgson score of c_i , denoted $\text{score}_E^D(c_i)$, is the smallest number of sequential swaps of adjacent candidates on the preference lists of voters in V needed to make c_i become the Condorcet winner. The candidates with the lowest score are the Dodgson election’s winners. That is, Dodgson defined his system to elect those candidates that are closest to being Condorcet winners in the sense of adjacent-swaps distance. Although Dodgson’s eighteenth-century election system was directly defined in terms of distance, there remains ongoing interest in understanding the classes of voting rules that can be captured in various distance-based frameworks (see, e.g., [MN08, EFS09]).

Unfortunately, it is known that deciding whether a given candidate is a winner according to Dodgson’s rule is quite complex. In fact, Hemaspaandra, Hemaspaandra, and Rothe [HHR97], strengthening an NP-hardness result of Bartholdi, Tovey, and Trick [BTT89b], showed that this problem is complete for parallelized access to NP. That is, it is complete for the Θ_2^P level of the polynomial hierarchy. Nonetheless, many researchers have sought efficient ways of computing Dodgson winners, for example by using

¹⁰Dodgson is better known as Lewis Carroll, the renowned author of “Alice’s Adventures in Wonderland.”

frequently correct heuristics [HH09, MPS08], fixed-parameter tractability (see [BTT89b, FHH09a, BGN10] and the discussion in Footnote 17 of [FHH09a]), and approximation algorithms for Dodgson scores [CCF⁺09].

In addition to its high computational cost in determining winners, Dodgson’s rule is often criticized for not having basic properties one would expect a good voting rule to have. For example, Dodgson’s rule is not “weakCondorcet consistent” (equivalently, it does not satisfy Fishburn’s “strict Condorcet principle”) [BBHH10] and doesn’t satisfy homogeneity and monotonicity (see [Bra09], which surveys a number of defects of Dodgson’s rule). We provide definitions for the latter two notions, as they will be relevant to this section.

Homogeneity. We say that a voting rule \mathcal{R} is homogeneous if for each election $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$, it holds that \mathcal{R} has the same winner set on E as on $E' = (C, V')$, where $V' = (v_1, v_1, v_2, v_2, \dots, v_n, v_n)$.

Monotonicity. We say that a voting rule \mathcal{R} is monotone if for each election $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$, it holds that if some candidate $c_i \in C$ is a winner of E then c_i is also a winner of an election E' that is identical to E except that some voters rank c_i higher (without changing the relative order of all the remaining candidates).

Continuing the Caragiannis et al. [CCF⁺09] line of research on approximately computing Dodgson scores, Caragiannis et al. [CKKP10] devised an approximation algorithm for computing Dodgson scores that, given an election $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$ and a candidate c_i in C , computes in polynomial time a nonnegative integer $\text{sc}_E(c_i)$ such that $\text{score}_E^D(c_i) \leq \text{sc}_E(c_i)$ and $\text{sc}_E(c_i) = O(m \log m) \cdot \text{score}_E^D(c_i)$. That is, the algorithm given by Caragiannis et al. [CKKP10] is, in a natural sense, an $O(m \log m)$ -approximation of the Dodgson score.¹¹ This algorithm has additional properties: If one defines a voting rule to elect those candidates that have lowest scores according to the algorithm, then that voting rule is Condorcet consistent (i.e., when a Condorcet winner exists, he or she is the one and only winner under the voting rule), homogeneous, and monotone.

The result of Caragiannis et al. [CKKP10] is very interesting, but unfortunately the voting rule defined by their approximation algorithm is somewhat complicated and arguably might seem not to be very natural. We now show that the maximin rule—which like the Caragiannis et al. rule is Condorcet-consistent, homogeneous, and monotone, but which in addition is a long-existing and natural rule—also elects candidates that are, in a certain different yet precise sense, “close” to being Dodgson winners. Our proof is inspired by that of Caragiannis et al. [CKKP10].

¹¹Throughout this section, we use the notion “ $f(m)$ -approximation of g ” in the sense it is typically used when dealing with minimization problems. That is, we mean that the approximation outputs a value that is at least g and at most $f(m) \cdot g$. We slightly abuse the interaction between this notation and Big-Oh notation, in the quite standard and intuitive way. And we assume that the argument domain that g and the approximation share is clear from context—in this paper, their arguments are an election E and a candidate c_i .

Theorem 4.7. Let $E = (C, V)$ be an election and let $W \subseteq C$ be a set of candidates that win in E according to the maximin rule. Let $m = \|C\|$ and let $s = \min_{c_i \in C} \text{score}_E^D(c_i)$. For each $c_i \in W$ it holds that $s \leq \text{score}_E^D(c_i) \leq m^2 s$.

Proof. Let us fix an election $E = (C, V)$ with $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$. For each two candidates $c_i, c_j \in C$ we define $\text{df}_E(c_i, c_j)$ to be the smallest number k such that if k voters in V changed their preference order to rank c_i ahead of c_j , then c_i would be preferred to c_j by more than half of the voters. Note that if for some $c_i, c_j \in C$ we have $\text{df}_E(c_i, c_j) > 0$ then

$$N_E(c_i, c_j) + \text{df}_E(c_i, c_j) = \left\lfloor \frac{n}{2} \right\rfloor + 1.$$

For each candidate $c_i \in C$ we define $\text{sc}'_E(c_i)$ to be

$$\text{sc}'_E(c_i) = m^2 \max\{\text{df}_E(c_i, c_j) \mid c_j \in C - \{c_i\}\}.$$

We now prove that sc' is an m^2 -approximation of the Dodgson score.

Lemma 4.8. For each $c_i \in C$ it holds that $\text{score}_E^D(c_i) \leq \text{sc}'_E(c_i) \leq m^2 \text{score}_E^D(c_i)$.

Proof. Let us fix some $c_i \in C$. To see that the second inequality in the lemma statement holds, note that $\max\{\text{df}_E(c_i, c_j) \mid c_j \in C - \{c_i\}\} \leq \sum_{c_j \in C - \{c_i\}} \text{df}_E(c_i, c_j) \leq \text{score}_E^D(c_i)$ because for each candidate c_k we, at least, have to perform $\text{df}_E(c_i, c_k)$ swaps to ensure that c_i defeats c_k in their majority head-to-head contest. Thus, after multiplying by m^2 , we have

$$m^2 \max\{\text{df}_E(c_i, c_j) \mid c_j \in C - \{c_i\}\} \leq m^2 \text{score}_E^D(c_i).$$

Let us now consider the first inequality. Let c_k be some candidate in $C - \{c_i\}$. To make sure that c_i is ranked higher than c_k by more than half of the voters, we can shift c_i to the first position in the preference lists of $\max\{\text{df}_E(c_i, c_j) \mid c_j \in C - \{c_i\}\} \geq \text{df}_E(c_i, c_k)$ voters (or, all the remaining voters if less than $\max\{\text{df}_E(c_i, c_j) \mid c_j \in C - \{c_i\}\}$ voters do not rank c_i as their top choice). This requires at most m adjacent swaps per voter. Since there are $m - 1$ candidates in $C - \{c_i\}$, $m^2 \max\{\text{df}_E(c_i, c_j) \mid c_j \in C - \{c_i\}\}$ adjacent swaps are certainly sufficient to make c_i a Condorcet winner. (Lemma 4.8) \square

It remains to show that if some candidate c_i is a maximin winner in E then $\text{sc}'_E(c_i)$ is minimal. Fortunately, this is easy to see. If some candidate c_i is a Condorcet winner of E then he or she is the unique maximin winner and he or she is the unique candidate c_i with $\text{sc}'_E(c_i) = 0$. Let us assume that there is no Condorcet winner of E . Let us fix some candidate $c_i \in C$ and let $c_k \in C - \{c_i\}$ be a candidate such that $\text{sc}'_E(c_i) = m^2 \text{df}_E(c_i, c_k)$. That is, $\text{df}_E(c_i, c_k) = \max\{\text{df}_E(c_i, c_j) \mid c_j \in C - \{c_i\}\}$ and $\text{df}_E(c_i, c_k) > 0$. Due to this last fact and our choice of c_k , we have $\text{df}_E(c_i, c_k) = \left\lfloor \frac{n}{2} \right\rfloor + 1 - N_E(c_i, c_k)$ and so

$$N_E(c_i, c_k) = \left\lfloor \frac{n}{2} \right\rfloor + 1 - \text{df}_E(c_i, c_k) = \min_{c_j \in C - \{c_i\}} N_E(c_i, c_j) = \text{score}_E(c_i),$$

where $\text{score}_E(c_i)$ is the maximin score of c_i in E . Thus each candidate c_i with the lowest value $\text{sc}'_E(c_i)$ also has the highest maximin score. \square

Theorem 4.7 says that every maximin winner’s Dodgson score is no less than the Dodgson score of the Dodgson winner(s) (that fact of course holds trivially), and is no more than m^2 times the Dodgson score of the Dodgson winner(s). That is, we have proven that *no candidate whose Dodgson score is more than m^2 times that of the Dodgson winner(s) can be a maximin winner.*

Since maximin is Condorcet consistent, homogeneous, and monotone, our result interestingly relates to the approximation of Caragiannis et al. [CKKP10], who achieved an $O(m \log m)$ -approximation of Dodgson score while maintaining Condorcet consistency, homogeneity and monotonicity (recall the discussion before Theorem 4.7). Admittedly, our “closeness” factor is m^2 , which is worse than achieving $O(m \log m)$. And our closeness is in a different sense, since our theorem is applying its bound just between Dodgson scores, and just on the winner set. In contrast, Caragiannis et al. [CKKP10] and even our own Lemma 4.8 relate the Dodgson score to the Caragiannis et al. score and the sc' score, and those approximations hold for all candidates. However, we achieve our m^2 closeness factor for a voting rule, maximin, that is well known and natural.

5 Fixed-Parameter Tractability

In this section we consider the parameterized complexity of multipronged control, in particular, the case where we can assume that the number of candidates is a small constant. Elections with few candidates are very natural: For example, in many countries presidential elections involve only a handful of candidates. The reader can easily imagine many other examples.

The main result of this section is that for many natural election systems \mathcal{E} (formally, for all election systems whose winner determination problem can be expressed via an integer linear program of a certain form), it holds that the $\mathcal{E}\text{-AC+DC+AV+BV+BV}$ control problem is fixed-parameter tractable (is in the complexity class FPT) for the parameter “number of candidates,” both in the constructive setting and in the destructive setting. This result combines and significantly enhances FPT results from the literature, in particular, from the papers [FHH09a, FHHR09a], which are the model for and inspiration of this section. We also make explicit an “automatic” path to such results that is implicit in [FHH09a, FHHR09a]. This path should be helpful in letting many future analyses be done as tool-application exercises, rather than being case-by-case challenges.

In this section we focus exclusively on the number of candidates as our parameter. That is, our parameter is the number of candidates initially in the election plus the number of candidates (if any) in the set of potential additional candidates. That is, in terms of the variables we have been using to describe multiprong control the parameter is $\|C\| + \|A\|$.

We mention that researchers sometimes analyze other parameterizations. For example, Liu et al. [LFZL09], Liu and Zhu [LZ10], and Betzler and Uhlmann [BU09] consider as the parameter the amount of change that one is allowed to use (e.g., the number of candidates one can add), Bartholdi, Tovey, and Trick [BTT89b], Betzler and Uhlmann [BU09], and Faliszewski et al. [FHHR09a] study as the parameter the number of voters (and also

sometimes the number of candidates). And other parameters are sometimes used when considering the so-called possible winner problem, see, e.g., [BD09,BHN09]. However, we view the parameter “number of candidates” as the most essential and the most natural one. We now proceed with our discussion of fixed-parameter tractability, with the number of candidates as the parameter.

Let us consider an election system \mathcal{E} and a set $C = \{c_1, \dots, c_m\}$ of candidates. There are exactly $m!$ preference orders over the candidates in C and we will refer to them as $o_1, \dots, o_{m!}$. Let us assume that \mathcal{E} is anonymous (i.e., the winners of each \mathcal{E} election do not depend on the order of votes or the names of the voters, but only—for each preference order o_i —on the number of votes with that preference order). We define predicate $win_{\mathcal{E}}(c_j, n_1, \dots, n_{m!})$ to be true if and only if c_j is a unique winner of \mathcal{E} elections with $C = \{c_1, \dots, c_m\}$, where for each i , $1 \leq i \leq m!$, there are exactly n_i voters with preference order o_i . For the rest of this section, our inequalities always use one of the four operators “ $>$,” “ \geq ,” “ $<$,” and “ \leq .¹²

Definition 5.1. *We say that an anonymous election system \mathcal{E} is unique-winner (nonunique-winner) integer-linear-program implementable if for each set of candidates $C = \{c_1, \dots, c_m\}$ and each candidate $c_j \in C$ there exists a set S of linear inequalities with variables $n_1, \dots, n_{m!}$ such that:*

1. *If the integer assignment $n_1 = \hat{n}_1, \dots, n_{m!} = \hat{n}_{m!}$ satisfies S , then each \hat{n}_i belongs to \mathbb{N} ,¹³*
2. *S can be computed (i.e., obtained) in time polynomial in $m!$,¹⁴ and*
3. *for each $(\hat{n}_1, \dots, \hat{n}_{m!}) \in \mathbb{N}^{m!}$, we have that (a) holds if and only if (b) holds, where (a) and (b) are as follows:*

$$(a) \quad S \text{ is satisfied by the assignment } n_1 = \hat{n}_1, \dots, n_{m!} = \hat{n}_{m!}.$$

¹²We allow both strict and nonstrict inequalities. Since we allow only integer solutions, it is easy to simulate strict inequalities with nonstrict ones and to simulate nonstrict inequalities with strict ones, in both cases simply by adding a “1” to the appropriate side of the inequality. So we could equally well have allowed just strict, or just nonstrict, inequalities.

¹³It is easy to put $m!$ inequalities into S enforcing this condition. And this condition will help us make the electoral part of our definition meaningful, i.e., it will avoid having problems from the restriction in the final part of this definition that lets us avoid discussing negative numbers of voters.

¹⁴We mention in passing that if the $m!$ in this part of the definition were changed to any other computable function of m , e.g., $m^{m^{m^m}}$, we would still obtain FPT results, and still would have them hold even in the strengthened version of FPT in which the f of “ $f(\text{parameter}) \cdot \text{Inputsize}^{O(1)}$ ” is required to be computable. However, due to $m!$ being the number of preference orders over m candidates, having S be obtainable in time polynomial in $m!$ will in practice be a particularly common case.

We also mention in passing that the FPT-establishing framework in this section and the results it yields, similarly to the case in our work mentioned earlier [FHH09a,FHHR09a], not only will apply in the model where votes are input as a list of ballots, one per person, but also will hold in the so-called “succinct” model (see [FHH09a,FHHR09a]), in which we are given the votes not as individual ballots but as binary numbers providing the number of voters having each preference order (or having each occurring preference order).

- (b) c_j is a unique winner (is a winner) of an \mathcal{E} election in which for each i , $1 \leq i \leq m!$, there are exactly \hat{n}_i voters with preference order o_i , where o_i is the i 'th preference order over the set C .

In a slight abuse of notation, for integer-linear-program implementable election systems \mathcal{E} we will simply refer to the set S of linear inequalities from Definition 5.1 as $win_{\mathcal{E}}(c_j, n_1, \dots, n_{m!})$. The particular set of candidates will always be clear from context. Naturally, it is easy to adapt Definition 5.1 to apply to approval voting, but for the sake of brevity we will not do so.

We are not aware of any natural systems that are integer-linear-program unique-winner implementable yet not integer-linear-program nonunique-winner implementable, or vice versa. In this paper we focus on the unique winner model so the reader may wonder why we defined the nonunique winner variant of integer-linear-program implementability. The answer is that, as we will see later in this section, it is a useful notion when dealing with destructive control.

The class of election systems that are integer-linear-program implementable is remarkably broad. For example, it is variously implicit in or a consequence of the results of [FHH09a] that plurality, veto, Borda, Dodgson, and each polynomial-time computable (in the number of candidates) family of scoring protocols are integer-linear-program implementable.¹⁵ For many other election systems (e.g., Kemeny [Kem59, YL78] and Copeland) it is not clear whether they are integer-linear-program implementable, but there are similar approaches that will be as useful for us. We will return to this issue at the end of this section.

Theorem 5.2. *Let \mathcal{E} be an integer-linear-program unique-winner implementable election system. For number of candidates as the parameter, constructive \mathcal{E} -AC+DC+AV+DV+BV is in FPT.*

Proof. Let $(C, A, V, W, p, k_{AC}, k_{DC}, k_{AV}, k_{DV}, k_{BV})$ be our input instance of the constructive \mathcal{E} -AC+DC+AV+DV+BV control problem, as described in Definition 3.1. Let $C = \{p, c_1, \dots, c_{m'}\}$ and $A = \{a_1, \dots, a_{m''}\}$. Our parameter, the total number of candidates, is $m = m' + m'' + 1$. For each subset K of $C \cup A$ we let $o_1^K, \dots, o_{\|K\|!}^K$ mean the $\|K\|!$ preference orders over K .

The idea of our algorithm is to perform an exhaustive search through all the subsets of candidates K , $K \subseteq C \cup A$, and for each K check whether (a) it is possible to obtain K from C by deleting at most k_{DC} candidates and adding at most k_{AC} candidates from A , and (b) it is possible to ensure that p is a unique winner of election (K, V) by deleting at most k_{DV} voters, adding at most k_{AV} voters from W , and bribing at most k_{BV} voters. Given K , step

¹⁵Let m be the number of candidates. A scoring protocol is a vector of m nonnegative integers satisfying $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_m$. Each candidate receives α_i points for each vote that ranks him or her in the i 'th position, and the candidate(s) with most points win. Many election systems can be viewed as families of scoring protocols. For example, plurality is defined by scoring protocols of the form $(1, 0, \dots, 0)$, veto is defined by scoring protocols of the form $(1, \dots, 1, 0)$, and Borda is defined by scoring protocols of the form $(m-1, m-2, \dots, 0)$, where m is the number of candidates.

(a) can easily be implemented in polynomial time. To implement step (b), we introduce a linear integer program $P(K)$, which is satisfiable if and only if step (b) holds. Let us now fix $K \subseteq C \cup A$ and describe the integer linear program $P(K)$.

We assume that $p \in K$ as it is not legal to delete p (and it would be pointless, given that we want to ensure his or her victory). We interpret preference orders of voters in V and W as limited to the candidate set K . We use the following *constants* in our program. For each i , $1 \leq i \leq \|K\|!$, we let n_i^V be the number of voters in V with preference order o_i^K , and we let n_i^W be the number of voters in W with preference order o_i^K . $P(K)$ contains the following variables (described together with their intended interpretation):

Variables $av_1, \dots, av_{\|K\|!}$. For each i , $1 \leq i \leq \|K\|!$, we interpret av_i as the number of voters with preference o_i^K that we add from W .

Variables $dv_1, \dots, dv_{\|K\|!}$. For each i , $1 \leq i \leq \|K\|!$, we interpret dv_i as the number of voters with preference o_i^K that we delete from V .

Variables $bv_{1,1}, bv_{1,2}, \dots, bv_{1,\|K\|!}, bv_{2,1}, \dots, bv_{\|K\|!,\|K\|!}$. For each i, j , $1 \leq i, j \leq \|K\|!$, we interpret $bv_{i,j}$ as the number of voters with preference o_i^K that, in case $i \neq j$, we bribe to switch to preference order o_j^K , or, in case $i = j$, we leave unbribed.

$P(K)$ contains the following constraints.

1. All the variables have nonnegative values.
2. For each variable av_i , $1 \leq i \leq \|K\|!$, there are enough voters in W with preference order o_i^K to be added. That is, for each i , $1 \leq i \leq \|K\|!$, we have a constraint $av_i \leq n_i^W$. Altogether, we can add at most k_{AV} voters so we have a constraint $\sum_{i=1}^{\|K\|!} av_i \leq k_{AV}$.
3. For each variable dv_i , $1 \leq i \leq \|K\|!$, there are enough voters in V with preference order o_i^K to be deleted. That is, for each i , $1 \leq i \leq \|K\|!$, we have a constraint $dv_i \leq n_i^V$. Altogether, we can delete at most k_{DV} voters so we have a constraint $\sum_{i=1}^{\|K\|!} dv_i \leq k_{DV}$.
4. For each variable $bv_{i,j}$, $1 \leq i, j \leq \|K\|!$, there are enough voters with preference o_i^K to be bribed. That is, for each i , $1 \leq i \leq \|K\|!$, we have a constraint $\sum_{j=1}^{\|K\|!} bv_{i,j} = n_i^V + av_i - dv_i$ (the equality comes from the fact that for each i , $1 \leq i \leq \|K\|!$, $bv_{i,i}$ is the number of voters with preference o_i^K that we do not bribe). Altogether, we can bribe at most k_{BV} voters so we also have a constraint

$$\left(\sum_{i=1}^{\|K\|!} \sum_{j=1}^{\|K\|!} bv_{i,j} \right) - \sum_{i=1}^{\|K\|!} bv_{i,i} \leq k_{BV}.$$

5. Candidate p is the unique winner of the election after we have executed all the adding, deleting, and bribing of voters. Using the fact that \mathcal{E} is integer-linear-program unique-winner implementable, we can express this as $\text{win}_{\mathcal{E}}(p, \ell_1, \dots, \ell_{\|K\|!})$, where we substitute each ℓ_j , $1 \leq j \leq \|K\|!$, by $\sum_{i=1}^{\|K\|!} bv_{i,j}$ (note that, by previous constraints, variables describing bribery already take into account adding and deleting voters). This is a legal integer-linear-program constraint as $\text{win}_{\mathcal{E}}(p, \ell_1, \dots, \ell_{\|K\|!})$ is simply a conjunction of linear inequalities over $\ell_1, \dots, \ell_{\|K\|!}$.

The number of variables and the number of inequalities in $P(K)$ are each polynomially bounded in $m!$. Keeping in mind Definitions 3.1 and 5.1, it is easy to see that program $P(K)$ does exactly what we expect it to. And testing whether $P(K)$ is satisfiable (i.e., has an integer solution, as we are in the framework of an integer linear program) is in FPT, with respect to the number of candidates being our parametrization, by using Lenstra's algorithm [Len83]. Thus our complete FPT algorithm for the \mathcal{E} -AC+DC+AV+DV+BV problem works as follows. For each subset K of $C \cup A$ that includes p we execute the following two steps:

1. Check whether it is possible to obtain K from C by deleting at most k_{DC} candidates and by adding at most k_{AC} candidates from A .
2. Form linear program $P(K)$ and check whether it has any integral solutions using the algorithm of Lenstra [Len83]. Accept if so.

If after trying all sets K we have not accepted, then reject.

From the previous discussion, this algorithm is correct. Also, since (a) there are exactly 2^{m-1} sets K to try, (b) executing the first step above can be done in time polynomial in m , and (c) the second step is in FPT (given that m is the parameter), constructive \mathcal{E} -AC+DC+AV+DV+BV is in FPT for parameter m . \square

The above theorem deals with constructive control only. However, using its proof, it is easy to prove a destructive variant of the result. We say that an election system is strongly voiced [HHR07] if it holds that whenever there is at least one candidate, there is at least one winner.

Corollary 5.3. *Let \mathcal{E} be a strongly voiced, integer-linear-program nonunique-winner implementable election system. Destructive \mathcal{E} -AC+DC+AV+DV+BV is in FPT for the parameter number of candidates.*

To see that the corollary holds, it is enough to note that for strongly voiced election systems a candidate can be prevented from being a unique winner if and only if some other candidate can be made a (possibly nonunique) winner (see e.g., Footnote 5 of [HHR07] for a relevant discussion). Thus to prove Corollary 5.3, we can simply use an algorithm that for each candidate other than the despised one sees whether that candidate can be made a (perhaps nonunique) winner, and if any can be made a (perhaps nonunique) winner, declares destructive control achievable. (And the precise integer linear programming

feasibility problem solution given by Lenstra's algorithm will reveal what action achieves the control.) This can be done in FPT using the algorithm from the proof of Theorem 5.2, adapted to work for the nonunique-winner problem (this is trivial given that Corollary 5.3 assumes that \mathcal{E} is integer-linear-program nonunique-winner implementable).

Let us now go back to the issue that some election systems may not be integer-linear-program implementable. As an example, let us consider maximin. Let $E = (C, V)$ be an election, where $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$. As before, by $o_1, \dots, o_{m!}$ we mean the $m!$ possible preference orders over C , and for each i , $1 \leq i \leq m!$, by n_i we mean the number of voters in V that report preference order o_i . For each c_i and c_j in C , $c_i \neq c_j$, we let $O(c_i, c_j)$ be the set of preference orders over C where c_i is preferred to c_j . Let $k = (k_1, \dots, k_m)$ be a vector of nonnegative integers such that for each i , $1 \leq i \leq m$, it holds that $1 \leq k_i \leq m$. For such a vector k and a candidate $c_\ell \in C$ we define $M(c_\ell, k_1, \dots, k_m)$ to be the following set of linear integer inequalities:

1. For each candidate c_i , his or her maximin score is equal to $N_E(c_i, c_{k_i})$. That is, for each i, j , $1 \leq i, j \leq m$, $i \neq j$, we have constraint $\sum_{o_k \in O(c_i, c_{k_i})} n_k \leq \sum_{o_k \in O(c_i, c_j)} n_k$
2. c_ℓ has the highest maximin score in election E and thus is the unique winner of E . That is, for each i , $1 \leq i \leq m$, $i \neq \ell$, we have constraint $\sum_{o_k \in O(c_\ell, c_{k_\ell})} n_k > \sum_{o_k \in O(c_i, c_{k_i})} n_k$.

It is easy to see that c_ℓ is a unique maximin winner of E if and only if there is a vector $k = (k_1, \dots, k_m)$ such that all inequalities of $M(c_\ell, k_1, \dots, k_m)$ are satisfied. It is also clear how to modify the above construction to handle the nonunique winner case. Since there are only $O(m^m)$ vectors k to try and each $M(c_\ell, k_1, \dots, k_m)$ contains $O(m^2)$ inequalities, it is easy to modify the proof of Theorem 5.2 to work for maximin: Assuming that one is interested in ensuring candidate c_ℓ 's victory, one simply has to replace program $P(K)$ in the proof of Theorem 5.2 with a family of programs that each include a different $M(c_\ell, k_1, \dots, k_m)$ for testing if c_ℓ had won. And one would accept if any of these were satisfiable. Thus we have the following result.

Corollary 5.4. *Constructive $AC+DC+AV+DV+BV$ control and destructive $AC+DC+AV+DV+BV$ control are both in FPT for maximin for the parameter number of candidates.*

The above construction for the winner problem in maximin can be viewed as, in effect, a disjunction of a set of integer linear programs. Such constructions for the winner problem have already been obtained for Kemeny in [FHH09a] and for Copeland in [FHHR09a]. Thus we have the following theorem.

Corollary 5.5. *With number of candidates as the parameter, constructive $AC+DC+AV+DV+BV$ control and destructive $AC+DC+AV+DV+BV$ control are in FPT for Kemeny and, for each each rational α , $0 \leq \alpha \leq 1$, for Copeland $^\alpha$.*

We conclude with an important caveat. The FPT algorithms of this section are very broad in their coverage, but in practice they would be difficult to use as their running time depends on (the fixed-value parameter) m in a very fast-growing way and as Lenstra’s algorithm has a large multiplicative constant in its polynomial running time. Thus the results of this section should best be interpreted as indicating that, for multipronged control in our setting, it is impossible to prove non-FPT-ness (and so it clearly is impossible to prove fixed-parameter *hardness* in terms of the levels of the so-called “W” hierarchy of fixed-parameter complexity, unless that hierarchy collapses to FPT). If one is interested in truly practically implementing a multipronged control attack, one should probably devise a problem-specific algorithm rather than using our very generally applicable FPT construction.

6 Conclusions

We have shown that combining various types of control into multiprong control attacks is a useful technique. It allows us to study more realistic control models, to express control vulnerability results and proofs in a compact way, and to obtain vulnerability results that are stronger than would be obtained for single prongs alone.

The main finding of our paper is that, to the extent to which we can draw conclusions from the set of election systems that we have studied, vulnerabilities to basic control types can often be combined to form a vulnerability to their multipronged control combination. (Table 2 summarizes our results regarding the five election systems we have focused on in this paper.) However, we have also seen that there exists a natural election system that is vulnerable to both constructive AC control and constructive AV control but that is resistant to constructive AC+AV control. We have also shown that as far as fixed-parameter tractability goes, at least with respect to the parameter number of candidates, a very broad class of election systems is vulnerable to the full AC+DC+AV+DV+BV control attack. And we have taken a small detour and proven that no candidate whose Dodgson score is more than $\|C\|^2$ times the Dodgson winner’s score can be a maximin winner.

This paper studies multipronged control where the prongs may include various standard types of control or bribery. However, it is easy to see that our framework can be naturally extended to include manipulation. To do so, one would have to allow some of the voters—the manipulators—to have blank preference orders and, if such voters were to be included in the election, the controlling agent would have to decide on how to fill them in. It is interesting that in this model the controlling agent might be able to add manipulative voters (if there were manipulators among the voters that can be added) or even choose to delete them (it may seem that deleting manipulators is never useful but Zuckerman, Procaccia, and Rosenschein [ZPR09] give an example where deleting a manipulator is necessary to make one’s favorite candidate a winner of a Copeland election).

We mention as a natural but involved open direction the study of multipronged control in the setting where there are multiple controlling agents, each with a different goal, each controlling a different prong. In such a setting, it is interesting to consider game-theoretic scenarios as well as situations in which, for example, one of the controlling agents is seeking

| Control type | plurality | | Condorcet | | Copeland ^{0.5} | | approval | | maximin | |
|-----------------|-----------|------|-----------|------|-------------------------|------|----------|------|---------|------|
| | Con. | Des. | Con. | Des. | Con. | Des. | Con. | Des. | Con. | Des. |
| AC | R | R | I | V | R | V | I | V | R | V |
| AC _u | R | R | I | V | R | V | I | V | V | V |
| DC | R | R | V | I | R | V | V | I | V | V |
| AV | V | V | R | V | R | R | R | V | R | R |
| DV | V | V | R | V | R | R | R | V | R | R |
| BV | V | V | R | V | R | R | R | V | R | R |

Table 2: Resistance to basic control types for the five main election systems studied in this paper. In the table, I means the system is immune to the given control type, R means resistance, and V means vulnerability. As shown in this paper, for each of the five election systems, all listed constructive vulnerabilities combine and all listed destructive vulnerabilities combine. All remaining prongs combine as described by Corollary 3.6. Constructive results for AC, AC_u, DC, AV, and DV for plurality and Condorcet are due to [BTT92] and their corresponding destructive results are due to [HHR07]. All results for AC ,AC_u, DC, AV, and DV for approval are due to [HHR07]. All results regarding Copeland, are due to [FHH09a]. Constructive bribery results for plurality and approval are due to [FHH09a], and the constructive bribery result for Condorcet is implicit in [FHH09a]. All the remaining results (i.e., all results regarding maximin, and destructive bribery results for plurality, approval, and Condorcet) are due to this paper.

an action that will succeed regardless of the action of the other attacker.

Acknowledgments

Supported in part by NSF grants CCF-0426761, IIS-0713061, and CCF-0915792, Polish Ministry of Science and Higher Education grant N-N206-378637, the Foundation for Polish Science’s Homing/Powroty program, AGH University of Science and Technology grant 11.11.120.865, the ESF’s EUROCORES program LogICCC, and Friedrich Wilhelm Bessel Research Awards to Edith Hemaspaandra and Lane A. Hemaspaandra. A preliminary version of this paper appeared in the proceedings of the 21st International Joint Conference on Artificial Intelligence, July 2009 [FHH09b]. We thank Edith Elkind and the anonymous IJCAI referees for helpful comments.

References

- [BBHH10] F. Brandt, M. Brill, E. Hemaspaandra, and L. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*. AAAI Press, July 2010. To appear.

- [BD09] N. Betzler and B. Dorn. Towards a dichotomy of finding possible winners in elections based on scoring rules. In *Proceedings of the 34th International Symposium on Mathematical Foundations of Computer Science*, pages 124–136. Springer-Verlag Lecture Notes in Computer Science #5734, August 2009.
- [BEH⁺] D. Baumeister, G. Erdélyi, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Computational aspects of approval voting. In J. Laslier and R. Sanver, editors, *Handbook of Approval Voting*. Springer. To appear.
- [BGN10] N. Betzler, J. Guo, and R. Niedermeier. Parameterized computational complexity of Dodgson and Young elections. *Information and Computation*, 208(2):165–177, 2010.
- [BHN09] N. Betzler, S. Hemmann, and R. Niedermeier. A multivariate complexity analysis of determining possible winners given incomplete votes. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 53–58. AAAI Press, July 2009.
- [Bla58] D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.
- [BO91] J. Bartholdi, III and J. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.
- [Bra09] F. Brandt. Some remarks on Dodgson’s voting rule. *Mathematical Logic Quarterly*, 55(4):460–463, 2009.
- [BTT89a] J. Bartholdi, III, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.
- [BTT89b] J. Bartholdi, III, C. Tovey, and M. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165, 1989.
- [BTT92] J. Bartholdi, III, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modeling*, 16(8/9):27–40, 1992.
- [BU09] N. Betzler and J. Uhlmann. Parameterized complexity of candidate control in elections and related digraph problems. *Theoretical Computer Science*, 410(52):43–53, 2009.
- [CCF⁺09] I. Caragiannis, J. Covey, M. Feldman, C. Homan, C. Kaklamanis, N. Karanikolas, A. Procaccia, and J. Rosenschein. On the approximability of Dodgson and Young elections. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1058–1067. Society for Industrial and Applied Mathematics, January 2009.

- [CKKP10] I. Caragiannis, C. Kaklamanis, N. Karanikolas, and A. Procaccia. Socially desirable approximations for Dodgson’s voting rule. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, pages 253–262. ACM Press, June 2010.
- [CS06] V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 627–634. AAAI Press, July 2006.
- [Dod76] C. Dodgson. A method of taking votes on more than two issues. Pamphlet printed by the Clarendon Press, Oxford, and headed “not yet published” (see the discussions in [MU95,Bla58], both of which reprint this paper), 1876.
- [DP08] S. Dobzinski and A. Procaccia. Frequent manipulability of elections: The case of two voters. In *Proceedings of the 4th International Workshop On Internet And Network Economics*, pages 653–664. Springer-Verlag Lecture Notes in Computer Science #5385, December 2008.
- [EFS09] E. Elkind, P. Faliszewski, and A. Slinko. On distance rationalizability of some voting rules. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*. ACM Press, July 2009.
- [EFS10] E. Elkind, P. Faliszewski, and A. Slinko. Cloning in elections. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*. AAAI Press, July 2010. To appear.
- [ENR09] G. Erdélyi, M. Nowak, and J. Rothe. Sincere-strategy preference-based approval voting fully resists constructive control and broadly resists destructive control. *Mathematical Logic Quarterly*, 55(4):425–443, 2009.
- [EPR10a] G. Erdélyi, L. Piras, and J. Rothe. Bucklin voting is broadly resistant to control. Technical Report arXiv:1005.4115 [cs.GT], arXiv.org, May 2010.
- [EPR10b] G. Erdélyi, L. Piras, and J. Rothe. Control complexity in fallback voting. Technical Report arXiv:1004.3398 [cs.GT], arXiv.org, April 2010.
- [FHH] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Using complexity to protect elections. *Communications of the ACM*. To appear.
- [FHH09a] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. How hard is bribery in elections? *Journal of Artificial Intelligence Research*, 35:485–532, 2009.
- [FHH09b] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Multimode attacks on elections. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 128–133. AAAI Press, July 2009.

- [FHHR09a] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research*, 35:275–341, 2009.
- [FHHR09b] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. A richer understanding of the complexity of election systems. In S. Ravi and S. Shukla, editors, *Fundamental Problems in Computing: Essays in Honor of Professor Daniel J. Rosenkrantz*, pages 375–406. Springer, 2009.
- [FHHR09c] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 118–127. ACM Press, July 2009.
- [FHS08] P. Faliszewski, E. Hemaspaandra, and H. Schnoor. Copeland voting: Ties matter. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, pages 983–990. International Foundation for Autonomous Agents and Multiagent Systems, May 2008.
- [FHS10] P. Faliszewski, E. Hemaspaandra, and H. Schnoor. Manipulation of Copeland elections. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, pages 367–374. International Foundation for Autonomous Agents and Multiagent Systems, May 2010.
- [FKN08] E. Friedgut, G. Kalai, and N. Nisan. Elections can be manipulated often. In *Proceedings of the 49th IEEE Symposium on Foundations of Computer Science*, pages 243–249. IEEE Computer Society, October 2008.
- [GJ79] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.
- [HH09] C. Homan and L. Hemaspaandra. Guarantees for the success frequency of an algorithm for finding Dodgson-election winners. *Journal of Heuristics*, 15(4):403–423, 2009.
- [HHR97] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Exact analysis of Dodgson elections: Lewis Carroll’s 1876 voting system is complete for parallel access to NP. *Journal of the ACM*, 44(6):806–825, 1997.
- [HHR07] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence*, 171(5–6):255–285, 2007.
- [HHR09] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Hybrid elections broaden complexity-theoretic resistance to control. *Mathematical Logic Quarterly*, 55(4):397–424, 2009.

- [HP01] G. Hägele and F. Pukelsheim. The electoral writings of Ramon Llull. *Studia Lulliana*, 41(97):3–38, 2001.
- [Kem59] J. Kemeny. Mathematics without numbers. *Daedalus*, 88:577–591, 1959.
- [Len83] H. Lenstra, Jr. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983.
- [LFZL09] H. Liu, H. Feng, D. Zhu, and J. Luan. Parameterized computational complexity of control problems in voting systems. *Theoretical Computer Science*, 410(27–29):2746–2753, 2009.
- [LZ10] H. Liu and D. Zhu. Parameterized complexity of control problems in maximin election. *Information Processing Letters*, 110(10):383–388, 2010.
- [Men10] C. Menton. Normalized range voting broadly resists control. Technical Report arXiv:1005.5698 [cs.GT], arXiv.org, May 2010.
- [ML06] I. McLean and H. Lorrey. Voting in the medieval papacy and religious orders. Report 2006-W12, Nuffield College Working Papers in Politics, Oxford, Great Britain, September 2006.
- [MLCM10] N. Maudet, J. Lang, Y. Chevaleyre, and J. Monnot. Possible winners when new candidates are added: The case of scoring rules. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, July 2010. To appear.
- [MN08] T. Meskanen and H. Nurmi. Closeness counts in social choice. In M. Braham and F. Steffen, editors, *Power, Freedom, and Voting*. Springer-Verlag, 2008.
- [MPRZ08] R. Meir, A. Procaccia, J. Rosenschein, and A. Zohar. The complexity of strategic behavior in multi-winner elections. *Journal of Artificial Intelligence Research*, 33:149–178, 2008.
- [MPS08] J. McCabe-Dansted, G. Pritchard, and A. Slinko. Approximability of Dodgson’s rule. *Social Choice and Welfare*, 31(2):311–330, 2008.
- [MU95] I. McLean and A. Urken. *Classics of Social Choice*. University of Michigan Press, 1995.
- [Nie06] R. Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006.
- [Pap94] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [Wal09] T. Walsh. Where are the really hard manipulation problems? The phase transition in manipulating the Veto rule. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 324–329. AAAI Press, July 2009.

- [XC08a] L. Xia and V. Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 109–118. ACM Press, July 2008.
- [XC08b] L. Xia and V. Conitzer. A sufficient condition for voting rules to be frequently manipulable. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 99–108. ACM Press, July 2008.
- [XCP09] L. Xia, V. Conitzer, A. Procaccia, and J. Rosenschein. Complexity of unweighted manipulation under some common voting rules. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 348–353. AAAI Press, July 2009.
- [YL78] H. Young and A. Levenglick. A consistent extension of Condorcet’s election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.
- [ZPR09] M. Zuckerman, A. Procaccia, and J. Rosenschein. Algorithms for the coalitional manipulation problem. *Artificial Intelligence*, 173(2):392–412, 2009.